

AD-A183 418

GENERATING NATURAL LANGUAGE UNDER PRAGMATIC CONSTRAINTS 1/3


(U) VALE UNIV NEW HAVEN CT DEPT OF COMPUTER SCIENCE


E H HOVY MAR 87 VALEU/CSD/RR-521 N00014-85-K-0108

UNCLASSIFIED

F/G 5/7

NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

AD-A183 418

7

DTIC
ELECTE
AUG 1 8 1987
S D D



**Generating Natural Language
Under Pragmatic Constraints**

Eduard Hendrik Hovy
YALEU/CSD/RR #521

March 1987

DISTRIBUTION STATEMENT A
Approved for public release
Distribution Unlimited

**YALE UNIVERSITY
DEPARTMENT OF COMPUTER SCIENCE**

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 521	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Generating Natural Language under Pragmatic Constraints		5. TYPE OF REPORT & PERIOD COVERED Research report
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Eduard Hendrik Hovy		8. CONTRACT OR GRANT NUMBER(s) N00014-85-K-0108
9. PERFORMING ORGANIZATION NAME AND ADDRESS Yale University Department of Computer Science 10 Hillhouse Avenue New Haven, CT 06520		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, VA 22209		12. REPORT DATE March 1987
		13. NUMBER OF PAGES 268
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Office of Naval Research Information Systems Program Arlington, VA 22217		15. SECURITY CLASS. (of this report) unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Artificial Intelligence Natural Language Processing Knowledge Representation Language Generation Cognitive Modelling		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) How and why do we say the same things in different ways to different people? Natural Language generation programs written to date are not sensitive to anything but their input, and therefore produce the same output to all hearers in all circumstances. For a generator to produce the various versions of a single input it must have the ability to make appropriate choices about the content and form of its text.		

Due to the flexibility of language, speakers can communicate far more than just the literal content of the words they use; the additional information usually serves some of their audience-related goals. Thus the generator's choices should be governed by the goals resulting from pragmatic issues such as its knowledge of the hearer and the setting of the conversation.

This research investigates the types of interpersonal goals that speakers can have; how they can be made specific enough to direct the generation process; the interpretation of input, but generator-directed inference, to find suitable forms of expression; and the representation of language in a phrasal lexicon. In the model of generation that incorporates these goals, planning and realization are interleaved processes, where the interleaving takes place at choice points. This view supports the standard top-down planning-to-realization approach as well as the bottom-up approach in which linguistic options present themselves as opportunities for the achievement of the active goals.

To illustrate these ideas, the program PAULINE (Planning And Uttering Language In Natural Environments) produces, from a single set of input representations in each of three domains, various paragraphs that differ in slant, content, and style.

OFFICIAL DISTRIBUTION LIST

Defense Documentation Center Cameron Station Alexandria, Virginia 22314	12 copies
Office of Naval Research Information Systems Program Code 437 Arlington, Virginia 22217	2 copies
Dr. Judith Daly Advanced Research Projects Agency Cybernetics Technology Office 1400 Wilson Boulevard Arlington, Virginia 22209	3 copies
Office of Naval Research Branch Office - Boston 495 Summer Street Boston, Massachusetts 02210	1 copy
Office of Naval Research Branch Office - Chicago 536 South Clark Street Chicago, Illinois 60615	1 copy
Office of Naval Research Branch Office - Pasadena 1030 East Green Street Pasadena, California 91106	1 copy
Mr. Steven Wong New York Area Office 715 Broadway - 5th Floor New York, New York 10003	1 copy
Naval Research Laboratory Technical Information Division Code 2627 Washington, D.C. 20375	6 copies
Dr. A.L. Slafkosky Commandant of the Marine Corps Code RD-1 Washington, D.C. 20380	1 copy
Office of Naval Research Code 455 Arlington, Virginia 22217	1 copy

Office of Naval Research Code 458 Arlington, Virginia 22217	1 copy
Naval Electronics Laboratory Center Advanced Software Technology Division Code 5200 San Diego, California 92152	1 copy
Mr. E.H. Gleissner Naval Ship Research and Development Computation and Mathematics Department Bethesda, Maryland 20084	1 copy
Captain Grace M. Hopper, USNR Naval Data Automation Command, Code 00H Washington Navy Yard Washington, D.C. 20374	1 copy
Dr. Robert Engelmores Advanced Research Project Agency Information Processing Techniques 1400 Wilson Boulevard Arlington, Virginia 22209	2 copies
Professor Omar Wing Columbia University in the City of New York Department of Electrical Engineering and Computer Science New York, New York 10027	1 copy
Office of Naval Research Assistant Chief for Technology Code 200 Arlington, Virginia 22217	1 copy
Computer Systems Management, Inc. 1300 Wilson Boulevard, Suite 102 Arlington, Virginia 22209	5 copies
Ms. Robin Dillard Naval Ocean Systems Center C2 Information Processing Branch (Code 8242) 271 Catalina Boulevard San Diego, California 92152	1 copy
Dr. William Woods BBN 50 Moulton Street Cambridge, MA 02138	1 copy

Professor Van Dam Dept. of Computer Science Brown University Providence, RI 02912	1 copy
Professor Eugene Charniak Dept. of Computer Science Brown University Providence, RI 02912	1 copy
Professor Robert Wilensky Univ. of California Elec. Engr. and Computer Science Berkeley, CA 94707	1 copy
Professor Allen Newell Dept. of Computer Science Carnegie-Mellon University Schenley Park Pittsburgh, PA 15213	1 copy
Professor David Waltz Univ. of Ill at Urbana-Champaign Coordinated Science Lab Urbana, IL 61801	1 copy
Professor Patrick Winston MIT 545 Technology Square Cambridge, MA 02139	1 copy
Professor Marvin Minsky MIT 545 Technology Square Cambridge, MA 02139	1 copy
Professor Negroponte MIT 545 Technology Square Cambridge, MA 02139	1 copy
Professor Jerome Feldman Univ. of Rochester Dept. of Computer Science Rochester, NY 14627	1 copy
Dr. Nils Nilsson Stanford Research Institute Menlo Park, CA 94025	1 copy

Dr. Alan Meyrowitz Office of Naval Research Code 437 800 N. Quincy Street Arlington, VA 22217	1 copy
LCOL Robert Simpson IPTO-DARPA 1400 Wilson Blvd Arlington, VA 22209	1 copy
Dr. Edward Shortliffe Stanford University MYCIN Project TC-117 Stanford Univ. Medical Center Stanford, CA 94305	1 copy
Dr. Douglas Lenat Stanford University Computer Science Department Stanford, CA 94305	1 copy
Dr. M.C. Harrison Courant Institute Mathematical Science New York University New York, NY 10012	1 copy
Dr. Morgan University of Pennsylvania Dept. of Computer Science & Info. Sci. Philadelphia, PA 19104	1 copy
Mr. Fred M. Griffie Technical Advisor C3 Division Marine Corps Development and Education Command Quantico, VA 22134	1 copy
Dr. Vince Sigilitto Program Manager AFOSR/NM Bolling Airforce Base Building 419 Washington, DC 20332	1 copy

**Generating Natural Language
Under Pragmatic Constraints**

Eduard Hendrik Hovy
YALEU/CSD/RR #521

March 1987

This work was supported by the Defense Advanced Research Projects Agency monitored by the Office of Naval Research under contract N00014-85-K-0108.

ABSTRACT

GENERATING NATURAL LANGUAGE UNDER PRAGMATIC CONSTRAINTS

Eduard Hendrik Hovy

Yale University

1987

How and why do we say the same things in different ways to different people?

Natural language generation programs written to date are not sensitive to anything but their input, and therefore produce the same output to all hearers in all circumstances. For a generator to produce various versions of a single input it must have the ability to make appropriate choices about the content and form of its text.

Due to the flexibility of language, speakers can communicate far more than just the literal content of the words they use; the additional information usually serves some of their audience-related goals. Thus the generator's choices should be governed by the goals resulting from pragmatic issues such as its knowledge of the hearer and the setting of the conversation.

This research investigates the types of interpersonal goals that speakers can have; how they can be made specific enough to direct the generation process; the interpretation of input, by generator-directed inference, to find suitable forms of expression; and the representation of language in a phrasal lexicon. In the model of generation that incorporates these goals, planning and realization are interleaved processes, where the interleaving takes place at choice points. This view supports the standard top-down planning-to-realization approach as well as the bottom-up approach in which linguistic options present themselves as opportunities for the achievement of the active goals.

To illustrate these ideas, the program PAULINE (Planning And Uttering Language In Natural Environments) produces, from a single set of input representations in each of three domains, various paragraphs that differ in slant, content, and style.

**Generating Natural Language
Under Pragmatic Constraints**

A Dissertation

Presented to the Faculty of the Graduate School

of

Yale University

in Candidacy for the Degree of

Doctor of Philosophy

by

Eduard Hendrik Hovy

May 1987

© Copyright by Eduard Hendrik Hovy 1987

ALL RIGHTS RESERVED

Acknowledgements

The temptation exists, when writing this page, to take the department phone list and delete some of the names you really cannot put faces to. I suspect this has been done. And the opposite temptation exists as well: thank you, dear committee, and goodbye.

A little too brusque, perhaps.

Oh well, then, switch the formality on high and the verbosity on high and the gratefulness on effusive and let the phrasal lexicon do its stuff...

IN THE FIRST PLACE, I WISH TO EXPRESS MY HEARTFELT GRATITUDE TO MY ADVISOR, PROFESSOR ROGER SCHANK. ACTING IN LOCO PARENTIS, AS HE WAS NOT INFREQUENTLY WONT TO INFORM US, TO HIS STUDENTS, HE DOES A TRULY EXCELLENT JOB OF FOSTERING THE CREATIVE APPROACH, THE WIDE VIEW, AND AN ENERGETIC AND CHALLENGING ENVIRONMENT.

Stop. Not good. Quit now, or at least try the verbosity on low and the formality on medium!

I WOULD ALSO LIKE TO THANK PROFESSORS DREW MCDERMOTT AND BOB ABELSON, THE OTHER TWO MEMBERS OF MY THESIS COMMITTEE. I APPRECIATE THEIR TAKING TIME TO READ THIS DISSERTATION AND PROVIDE ME WITH COMMENTS. AND I APPRECIATE THE HELP AND COMMENTS GIVEN ME BY DR CHRIS RIESBECK.

Better? Nyaaah... Need less gratefulness and more force.

FOR DISCUSSION AND ANALYSIS AND OCCASIONAL HEALTHY DOSES OF CYNICISM, THERE'S NOBODY TO BEAT LARRY BIRNBAUM. THANKS, LB.

FOR WORKING ON PAULINE AND ITS PREDECESSORS AND PROVIDING IDEAS AND LANGUAGE SKILLS, THANKS TO GITA ASHOK, YANG-DONG LEE, ASHWIN RAM, AND JEFF GROSSMAN. THANKS TO ROD MCGUIRE FOR THE VERY FIRST PUSH AND TO STEVE LYTINEN AND NATALIE DEHN. THANKS TO MICHAEL FACTOR FOR TALKS AND COMMENTING PAPERS AND TO DAVID LITTLEBOY FOR MORE CYNICISM (HIS PECULIAR FLAVOR, THOUGH) AND FOR CHESS AND TO CHRIS OWENS FOR OCCASIONAL DISCUSSIONS. THANKS VERY MUCH TO BILL BAIN FOR JUDGE AND FOR DISCUSSIONS IN WHICH I LEARNED A LOT. THANKS TO PATTI AND DONNA AND MICHELLA AND PAULA AND RIA. AND TO THE T PEOPLE AND THE FACILITY STAFF. AND OF COURSE TO THE ADVANCED RESEARCH PROJECTS AGENCY, MONITORED BY THE OFFICE OF NAVAL RESEARCH, WHOSE CONTRACT N00014-82-K-0149 SUPPORTED MOST OF THIS WORK.

Well, *that's* not too interesting any more. Too much thanks. Cut it short. Haste on high.

THANKS TO ALL THE OTHER PEOPLE IN THE LAB AND OUTSIDE IT WHO MADE LIFE BEARABLE WHEN I GOT HOMESICK, ESPECIALLY TO ROB AND ILANA FOR CHAMBER MUSIC AND MOTOR MECHANICS; TO GABRIEL (ON AND OFF THE SQUASH COURT) AND PAOLA (INSIDE AND OUTSIDE THE ICE-CREAM STORE). VERY ESPECIALLY TO BEVERLY. (NOW I EXPECT A RECOGNITION IN YOUR ACKNOWLEDGEMENTS TOO.)

Final tally: syrupiness unhealthily high. Rewrite? No, that's a program's job.

Hmmm...

But then, so is writing in the first place...

Contents

1	Language: A Multi-Purpose Tool	1
1.1	Introduction	1
1.2	Some Generator Texts	6
1.3	A Short Overview	13
1.4	Conclusion	15
1.5	The Form of this Dissertation	16
2	Pragmatics	18
2.1	What is Pragmatics? — Some Definitions	18
2.2	Some Pragmatic Aspects of Conversations	21
2.2.1	Interlocutors' Factual Knowledge	24
2.2.2	Interlocutors' Opinions	26
2.2.3	Interlocutors' Emotional States	27
2.2.4	Interlocutors' Goals: Altering Future Actions	27

2.2.5	Interpersonal Relationships	28
2.2.6	Conversational Atmosphere	31
2.3	PAULINE's Pragmatic Settings and Goals	32
2.3.1	From Pragmatics to the Generation Process	34
2.4	Rhetorical Goals of Opinion	38
2.5	Rhetorical Goals of Style	43
2.6	Conclusion	46
3	Interpretation	48
3.1	The Problem	48
3.2	Interpretation in Generation	51
3.3	Controlling Inferences	54
3.3.1	Unfortunate Practical Realities	54
3.3.2	Bottom-Up Inference	56
3.3.3	Top-Down Inference	57
3.4	How PAULINE Does It	60
3.5	Determining the Appropriate Level of Detail	63
3.6	The Inclusion of Reminders	66
3.6.1	Reminders	66
3.6.2	Adding Color to the Text	68
3.7	Conclusion	70

4 Affect in Text	71
4.1 Introduction	71
4.2 Computing Affect	75
4.3 The Affect Rule and its Application	76
4.4 Content: Topic Collection Techniques	78
4.4.1 Evasion	79
4.4.2 Selectivity	81
4.5 Form: Generating with Affect	85
4.5.1 Topic Organization	85
4.5.2 Sentence Inclusion	90
4.5.3 Sentence Organization	93
4.5.4 Clause Content and Organization	95
4.5.5 Word Choice	96
4.6 When and Where to be Partial	99
4.6.1 Timidity: The Inclusion of Sensitive Topics	99
4.6.2 The Degree of Partiality	101
4.7 Conclusion	103

5	Creating Style	107
5.1	The Nature of Style	107
5.2	Formality	108
5.3	Haste	117
5.4	Force	124
5.5	Summary	130
5.5.1	Topic Collection	133
5.5.2	Topic Organization	134
5.5.3	Sentence Organization	135
5.5.4	Analysis	135
5.6	Conclusion	138
6	Grammar and a Phrasal Lexicon	139
6.1	The Three Tasks of a Generator	139
6.2	Formative Information	143
6.2.1	Rules of Grammar	143
6.2.2	Lexical Entries	147
6.2.3	Becker's Phrasal Lexicon	148
6.3	Syntax Specialists	151
6.3.1	PAULINE	151

6.3.2	Specialists and Phrase Structure Symbols	153
6.3.3	Relations among Elements of the Lexicon	155
6.4	A Phrasal Grammar	158
6.5	A Small Example of Realization	160
6.5.1	Input Topic	161
6.5.2	Sentence Content and Organization	162
6.5.3	Clause Organization	164
6.5.4	Word Choice	166
6.5.5	The Process Repeats	170
6.6	Conclusion	171
7	Planning and Realization	173
7.1	The Trouble with Traditional Planning	173
7.2	Limited-Commitment Planning	176
7.3	Performing Restrictive Planning: Monitoring	178
7.4	Planning in PAULINE	181
7.4.1	Program Architecture	181
7.4.2	Computing Affect	182
7.4.3	Rhetorical Goals of Opinion	186
7.4.4	Topic Collection	189

7.4.5	Topic Interpretation	190
7.4.6	Additional Topic Inclusion	193
7.4.7	Topic Juxtaposition	194
7.4.8	Topic Ordering	195
7.4.9	Syntax	195
7.4.10	The Planning Process, Step by Step	197
7.5	Conclusion	199
8	A Review of Generation Literature	201
8.1	The Three Questions of Generation	201
8.1.1	How shall I say it?	203
8.1.2	What shall I say?	206
8.1.3	Why should I say it?	209
8.2	Conclusion: Doing it Better Next Time	210
	Appendices	212
	Appendix A	212
	Appendix B	235

Chapter 1

Language: A Multi-Purpose Tool

1.1 Introduction

When you compare the language produced by people to the text produced by existing language generation programs, one thing becomes immediately clear: people can say the same thing in various ways to achieve various effects, and generators cannot. The generator described here, PAULINE (Planning And Uttering Language In Natural Environments), addresses this shortcoming.

It is straightforward to write a language generation program that produces impressive text by associating a sentence template (or some equivalent general grammatical form) with each representational item and then using a grammar to realize the template into surface form. Such a program, however, is not sensitive to anything but the input items, and therefore produces the same output to all hearers in all circumstances.

When we produce language, we *tailor out text to the hearer and to the situation*. This enables us to communicate more information than is contained in the literal meanings of our words; indeed, the additional information often has a stronger effect on the hearer than the literal content has. This information is carried by both the content and the form of the

text. As speakers and hearers, we attach various interpretations of the speaker, his goals, the hearer, and the conversational circumstances, to the various ways of expressing a single underlying representation. These interpretations are governed by rules. For example, we use such rules to understand the different meanings of each of the following sentences:

"Old Bill finally kicked the bucket last night!"

"We are not going to see Uncle Bill any more..."

"I am very sorry to have to tell you that Bill passed away"

Speakers use the rules to determine how to say what they want to say. In order to exhibit the same degree of flexibility of expression, generator programs require such rules too.

What, then, is the additional information that speakers can convey? Consider the different points of view the speaker communicates in each of the following four descriptions of an event that happened at Yale University in April 1986:

(a) *On April 4, concerned Yale students constructed a shantytown on Boesak Plaza as a reminder to those in Woodbridge Hall (and all over campus and the community) that Yale is complicit (sic) with the system of apartheid that creates shantytowns where thousands of blacks are forced to live in squalor and fear. The shantytown, Winnie Mandela City, served as a focal point of education concerning South Africa and Yale's investments there. At 5:30 am on April 14 the Yale Administration had the shantytown torn down and had 76 students and community members who were defending the shanties arrested. After a huge outcry, the Administration allowed the shanties to be rebuilt. We will not be silenced; we will continue to challenge the University on their moral failure. (From: protester literature; the protesters renamed the plaza after the South African churchman Allan Boesak)*

(b) *On April 4, a small group of students took over Beinecke Plaza and built some shanties; they wanted to force Yale to sell its stocks in companies with branches located in South Africa. The university asked the students to move*

the shanties to another location, but the students refused. The university then granted them permission to occupy the plaza until the end of the week, so that they could be there to be seen by the university's trustees, the Yale Corporation, at their meeting. But even after the meeting, the students refused to leave the plaza, and police had to clear the shanties. Later, the university relented, and gave them permission to rebuild the shanties. It also announced that it would send a fact-finding mission to South Africa. (Speaker: anti-divestment Yale student)

*(c) On April 4, students at Yale built a symbolic shantytown to protest their school's investments in companies doing business in South Africa. The college ordered the shanties destroyed. The police arrested 76 protesters when the shantytown was torn down. Local politicians and more than 100 faculty members criticized the action. A week after it had ordered the removal of the shantytown - named Winnie Mandela City, after the South African foe of apartheid - the shantytown was reconstructed and the administration agreed to allow it to remain standing. Concurrently, Yale announced that its trustees, the Yale Corporation, would soon send a fact-finding mission to South Africa to investigate the actions of corporations in which it owns between \$350 million and \$400 million of stock. (From: *New York Times*, Sunday, 27 April, Connecticut section)*

(d) Some students erected a shantytown to protest Yale's investments in companies that have operations in South Africa. The University tore it down and arrested several of them. The students continued to demonstrate and finally the university said they could put up the shantytown again. The university said it would investigate its investments in South Africa. (Speaker: neutral student)

Clearly the first two speakers incorporate strongly their opinions about the shantytown issue; the second two speakers seem more neutral but differ in level of formality. But how do you "incorporate opinions" and what does it mean to "seem more neutral" and to "be

formal"? There is no single item in the texts that can be pinpointed as carrying the opinion or setting the level or formality; rather, each text seems to contain a number of little clues, and these clues cumulatively convey a certain impression to the reader. What are these little clues? Where do they appear in language and how do we decide to use them? How do they interact? What other impressions — information such as the speaker's emotional state, his social status relative to the hearer, the ways he would like to influence the hearer's future behavior — can be incorporated into language?

Some additional, rather overt, techniques are visible in the following two example texts. The first is an excerpt from an open letter to Yale President Giamatti from the university's clerical and technical workers' labor union negotiating committee, November 9, 1984:

(e) It is time, in the best interests of all concerned, to settle the strike. It is our understanding that the University administration, as well as the Union, has received a document entitled 'A Statement of Purpose by the Coalition to End the Strike'. We appreciate the spirit of the document. Clearly, the community earnestly desires and needs a settlement, so that Yale can get back to what it is supposed to be. Our members earnestly want a settlement. You have said that you do, too, and we are prepared to take you at your word. We are willing to compromise significantly to achieve a settlement. Therefore, we propose the immediate resumption of negotiations on a daily basis.

In contrast, compare this excerpt from an open letter from president Giamatti to the Yale community, September 26, 1984:

(f) I write with great disappointment following Local 34's action in calling a strike against the University. The University negotiating team has made concerted efforts, lately with the help of the Mediator, Eva Robins, to find common ground and to bring about a fair and reasonable settlement of the outstanding issues in a manner satisfactory to both parties. But the agreement has not been achieved.

In both cases, the authors seem honestly to want a solution. But consider how the former is a rousing call to action, while the latter is reasoned and calm, and see how each side imputes blame to the other... The union "earnestly desires and needs a settlement"; they "are willing to compromise"; they will "take [Yale] at [its] word" that Yale wants a settlement too. Clearly Yale does not *really* want a settlement! Clearly the union has to prod an unwilling Yale into negotiating! In addition to being slanted, this letter is forceful, exhortatory; taken at face value, it obviously tries to influence Giamatti's future actions. Compare that with Giamatti's much calmer, more reasoned response, in which he writes "with great disappointment"; his team has "made concerted efforts"; the agreement "has not been achieved". It is clear that *he* is not to blame! In fact, his disappointment casts Giamatti as an optimist, a reasonable man who hopes others are reasonable too. The fact that he does not place the blame for the failure of the settlement explicitly onto the other party (say, with "But the Union did not want to come to an agreement"), underscores this impression. Clearly it is the union that is recalcitrant! Similar techniques are used in the shantytown texts. In example (a), for instance, the protesters use such phrases as "concerned Yale students"; "constructed a shantytown... as a reminder"; Yale "had 76 students... arrested"; "a huge outcry". Clearly, the protesters were well-meaning, harmless people with much popular support. But then what about the opponent's account (b) that contains phrases such as "a small group of students"; "took over Beinecke Plaza"; "they wanted to force Yale to..."; "police had to clear the shanties"? Obviously the university did its best to remain conciliatory toward a few radicalized, rapacious students with an axe to grind!

Techniques such as implying reasonableness, willingness to compromise, trust in the goodwill of the opponent, etc., are obviously part of the art of creating opinion-bearing and manipulative text. So are techniques such as maximizing or minimizing support, implying the use of force, and the use of appropriate adjectives and adverbs.

1.2 Some Generator Texts

PAULINE, the computer program described here, uses strategies based on these techniques to produce various texts from underlying representations. In all, it has been tried on three distinct episodes. The first set of examples are generated from a representation of the shantytown episode. From a single representation — a network of about 120 representation elements — PAULINE produces over 100 different texts. For example, as a quick informal description of the issue, PAULINE says:

Example 1.

YALE UNIVERSITY PUNISHED A NUMBER OF STUDENTS FOR BUILDING A SHANTYTOWN ON BEINECKE PLAZA BY ARRESTING 76 STUDENTS AND TEARING IT DOWN ONE MORNING IN EARLY APRIL. THE STUDENTS WANTED YALE TO DIVEST FROM COMPANIES DOING BUSINESS IN SOUTH AFRICA. FINALLY, THE UNIVERSITY GAVE IN AND ALLOWED THE STUDENTS TO REBUILD IT.

This is the kind of description one may hear from a passerby who knows about the issue. Compare this with the following version, in which PAULINE, now being a journalist, uses much more formal language for its newspaper article (including, for example, such conjunctive phrases as “so that” and “concurrently”, such verbs as “construct”, “request”, and “give permission”, and the passive mode):

Example 2.

IN EARLY APRIL, A SHANTYTOWN --- NAMED WINNIE MANDELA CITY --- WAS CONSTRUCTED BY SEVERAL STUDENTS ON BEINECKE PLAZA SO THAT YALE WOULD DIVEST FROM COMPANIES DOING BUSINESS IN SOUTH AFRICA. THE LOCAL COMMUNITY SUPPORTED THE ACTION. AT 5:30 AM ON APRIL 14, THE SHANTYTOWN WAS DESTROYED BY OFFICIALS; ALSO, AT THAT TIME, THE POLICE ARRESTED 76 STUDENTS. THE STUDENTS REQUESTED THAT THE UNIVERSITY GIVE THEM PERMISSION TO REASSEMBLE THE SHANTYTOWN

WHILE SEVERAL LOCAL POLITICIANS AND FACULTY MEMBERS EXPRESSED CRITICISM OF YALE'S ACTION. FINALLY, THE UNIVERSITY PERMITTED THE STUDENTS TO RECONSTRUCT IT AND, CONCURRENTLY, YALE ANNOUNCED THAT A COMMISSION WOULD GO TO SOUTH AFRICA IN JULY TO EXAMINE THE SYSTEM OF APARTHEID.

Examples 1 and 2 are fine ways to introduce someone to the topic. They are the kind of thing we would say if we were unbiased informers, such as passersby and reporters. But what if we were protesters and were trying to convince someone of our opinion? When PAULINE is set up to produce text supporting the protesters' cause, i.e., when it is given the goals to slant its description of the episode toward the protesters, it produces:

Example 3.

AS A REMINDER TO YALE UNIVERSITY TO DIVEST FROM COMPANIES DOING BUSINESS IN SOUTH AFRICA, A LARGE NUMBER OF CONCERNED STUDENTS ERECTED A SHANTYTOWN --- NAMED WINNIE MANDELA CITY --- ON BEINECKE PLAZA IN EARLY APRIL. THE LOCAL COMMUNITY EXPRESSED SUPPORT FOR THE STUDENTS' ACTION. THE UNIVERSITY TOLD THE STUDENTS TO ERECT THE SHANTYTOWN ELSEWHERE. LATER, AT 5:30 AM ON APRIL 14, THE SHANTYTOWN WAS DESTROYED BY OFFICIALS; ALSO, AT THAT TIME, THE POLICE ARRESTED 76 STUDENTS. THE STUDENTS REQUESTED THAT YALE UNIVERSITY GIVE THEM PERMISSION TO REASSEMBLE IT ON BEINECKE PLAZA AND AT THE SAME TIME SEVERAL LOCAL POLITICIANS AND FACULTY MEMBERS EXPRESSED CRITICISM OF YALE'S ACTIONS. FINALLY, YALE PERMITTED THE STUDENTS TO RECONSTRUCT THE SHANTYTOWN AND, CONCURRENTLY, THE UNIVERSITY ANNOUNCED THAT A COMMISSION WOULD GO TO SOUTH AFRICA IN JULY TO EXAMINE THE SYSTEM OF APARTHEID.

In contrast, when PAULINE has to produce the official Yale version, it says:

Example 4.

IN EARLY APRIL, A SMALL NUMBER OF STUDENTS WERE INVOLVED IN A CONFRONTATION WITH YALE UNIVERSITY OVER YALE'S INVESTMENT IN COMPANIES DOING BUSINESS IN SOUTH AFRICA. THE STUDENTS CONSTRUCTED A SHANTYTOWN --- NAMED WINNIE MANDELA CITY --- ON BEINECKE PLAZA IN ORDER TO FORCE THE UNIVERSITY TO DIVEST FROM THOSE COMPANIES. YALE REQUESTED THAT THE STUDENTS ERECT IT ELSEWHERE, BUT THEY REFUSED TO LEAVE. THE UNIVERSITY INTENDED TO BE REASONABLE. THE UNIVERSITY GAVE IT PERMISSION TO EXIST UNTIL THE MEETING OF THE YALE CORPORATION, BUT EVEN AFTER THAT THE STUDENTS STILL REFUSED TO MOVE. AT 5:30 AM ON APRIL 14, OFFICIALS HAD TO DISASSEMBLE THE SHANTYTOWN. FINALLY, YALE, BEING CONCILIATORY TOWARD THE STUDENTS, NOT ONLY PERMITTED THEM TO RECONSTRUCT IT, BUT ALSO ANNOUNCED THAT A COMMISSION WOULD GO TO SOUTH AFRICA IN JULY TO EXAMINE THE SYSTEM OF APARTHEID.

Of course, this is not the way a protester or Yale official would tell you the story in person. For one thing, they may cut it a little shorter, and for another, they will be less formal and more openly opinionated. PAULINE illustrates:

Example 5.

I AM ANGRY ABOUT YALE'S ACTIONS. THE UNIVERSITY HAD OFFICIALS DESTROY A SHANTYTOWN CALLED WINNIE MANDELA CITY ON BEINECKE PLAZA AT 5:30 AM ON APRIL 14. A LOT OF CONCERNED STUDENTS BUILT IT IN EARLY APRIL. NOT ONLY DID YALE HAVE OFFICIALS DESTROY IT, BUT THE POLICE ARRESTED 76 STUDENTS. AFTER THE LOCAL COMMUNITY'S HUGE OUTCRY, THE UNIVERSITY ALLOWED THE STUDENTS TO PUT THE SHANTYTOWN UP THERE AGAIN.

Example 6.

IT PISSES ME OFF THAT A FEW SHIFTLESS STUDENTS WERE OUT TO MAKE TROUBLE ON BEINECKE PLAZA ONE DAY: THEY BUILT A SHANTYTOWN, WINNIE MANDELA CITY, BECAUSE THEY WANTED YALE UNIVERSITY TO PULL THEIR MONEY OUT OF COMPANIES WITH BUSINESS IN SOUTH AFRICA. I AM HAPPY THAT OFFICIALS REMOVED THE SHANTYTOWN ONE MORNING. FINALLY, YALE GAVE IN AND LET THE SHITHEADS PUT IT UP AGAIN, AND YALE SAID THAT A COMMISSION WOULD GO TO SOUTH AFRICA TO CHECK OUT THE SYSTEM OF APARTHEID.

The second set of examples PAULINE can talk about concerns a fictitious primary election between politicians Carter and Kennedy during their race for the Democratic nomination for Presidential candidate. The numbers and dates are not to be taken seriously. Again (from a single network of about 50 representation elements), PAULINE can simply describe the outcome:

Example 7.

ON 20 FEBRUARY, CARTER AND KENNEDY WERE THE CANDIDATES IN A PRIMARY IN MICHIGAN. CARTER LOST TO KENNEDY BY 1335 VOTES. AT PRESENT, KENNEDY HAS A BETTER CHANCE OF GETTING THE NOMINATION THAN HE HAD BEFORE. CARTER IS ALSO CLOSER TO GETTING THE NOMINATION THAN HE WAS BEFORE. BOTH CARTER AND KENNEDY WANT TO GET THE NOMINATION.

However, this text is clearly inappropriate if the hearer already knows something about the nature of candidates, nominations, and primaries. If, say, the interlocutors are in a hurry -- for example, while making a long-distance phone call -- the following is better:

Example 8.

WELL, SO CARTER LOST THE PRIMARY TO KENNEDY BY 1335 VOTES.

Of course, if we had any feelings about the matter, we would not say it this way. Our biases sneak into what we say so easily and so often that producing slanted text seems much more natural than being neutral! Given the goal to support Kennedy, PAULINE says:

Example 9.

KENNEDY NARROWED CARTER'S LEAD BY GETTING ALL OF 21850 VOTES IN THE PRIMARY IN MICHIGAN. IN A SIMILAR CASE, CARTER DECREASED UDALL'S LEAD IN A PRIMARY IN 1976, AND HE EASILY TROUNCED UDALL TO BE NOMINATED BY 2600 DELEGATES. I AM GLAD THAT PRESENTLY KENNEDY IS CLOSER TO GETTING THE NOMINATION THAN BEFORE.

This example is clearly spoken by a Kennedy supporter: it focuses on Kennedy's victory and current standing, Carter's loss, and the way in which a front-runner can be overtaken. Note that Kennedy is still behind Carter so the most the program can claim is that Kennedy has *many* committed delegates. What happens if PAULINE is, instead, a Carter supporter? Of course, it should make the most of the fact that Carter is still ahead in committed delegates, as well as downplay his loss:

Example 10.

CARTER HAS GOT MANY DELEGATES AT PRESENT; WHAT'S MORE, HE HAS GOT MANY DELEGATES MORE THAN IN THE PAST. I THINK IT'S GREAT THAT HE HAS GOT MANY MORE DELEGATES THAN KENNEDY.

Of course, while slanting the text, the program can still vary its style. In contrast to example 9, but with the same sympathies, PAULINE can be formal (say, while making a speech):

Example 11.

I AM PLEASED TO INFORM YOU THAT CARTER HAS IMPROVED HIS CHANCES OF WINNING THE NOMINATION. AT THE PRESENT TIME, CARTER HAS MANY

MORE DELEGATES THAN HE HAD IN THE PAST; ALSO, CARTER HAS MANY MORE THAN KENNEDY DOES.

Here, PAULINE skirts the central issue, Carter's loss, concentrating on more pleasant aspects. But what would happen in an extreme case? what if you, a Carter supporter, are speaking to your boss, an irascible Kennedy man, under pressure? How would you even bring the topic up? Under these circumstances PAULINE chose to do:

Example 12.

...exactly nothing! (This occurred as a surprise. What happened was the following: Initially, the program didn't find any topics that it liked and the hearer also liked, and it didn't have time to search the story representation for other topics or to find ways of mitigating the unpleasant ones). The program can also be aggressive, if it is speaking to someone it dislikes and wants to anger:

Example 13.

BUSTER, I THINK IT'S WONDERFUL THAT YOUR HERO, THAT JERK CARTER, REALLY GOT JUST 20515 VOTES. SO CARTER GOT AS MANY AS 1335 VOTES FEWER THAN THAT WONDERFUL KENNEDY DID IN THE PRIMARY. FORTUNATELY, KENNEDY SUCCEEDED IN GETTING ALL OF 21850 VOTES; ALSO, HE HAS GOT MANY DELEGATES AT THE PRESENT TIME. THAT MEANS THAT HE HAS GOT SEVERAL DELEGATES MORE THAN HE DID IN THE PAST.

These are some of the variations that PAULINE can produce, from a single underlying representation, under different circumstance settings.

The third set of examples derives as output from the JUDGE program ([Bain 85, 86]) that models the sentencing behavior of a judge. In a neutral setting, when just describing a typical case, PAULINE says:

Example 14.

AFTER JIM BUMPED MIKE ONCE, THEY FOUGHT, AND EVENTUALLY HE WAS
KNOCKED TO THE GROUND BY MIKE. HE SLAPPED HIM A FEW TIMES MIKE
STABBED JIM, AND JIM DIED.

The program is given a lot more information about each action of the fight, and is able to select appropriate topics, and group them together. Being concise, it can say:

Example 15.

JIM DIED IN A FIGHT.

Example 16.

AFTER JIM BUMPED MIKE ONCE, THEY FOUGHT, AND EVENTUALLY MIKE
KILLED JIM.

Being more expansive, PAULINE can include information that the JUDGE program inferred and used to produce a sentence, such as the actors' intents and justifications, in order to slant its description toward the victim or the defendant. For example, when defending Mike, PAULINE says:

Example 17.

FIRST, THAT JERK JIM BUMPED MIKE HARD AND HURT HIM. THEN MIKE
JUST TAPPED JIM ONCE. 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.

etc...

MIKE REALIZED THAT JIM REALLY HURT HIM; ALSO, HE EXPECTED THAT
JIM'S GOAL WAS TO HURT HIM. JIM WAS STABBED. AS A RESULT, JIM
ONLY DIED.

(not taking Jim's death too seriously, there is ...). And, in contrast, when defending Jim, different aspects are selected from the same representation and the program says:

Example 18.

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. 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.

etc. .

MIKE COULD EXPECT THAT IF HE STABBED JIM SEVERAL TIMES HE WOULD KILL HIM. HE STABBED HIM. THE RESULT WAS THAT HE WOUNDED HIM. MIKE'S CAUSING JIM TO BE KILLED WAS NOT JUSTIFIED. AS A RESULT, JIM DIED.

1.3 A Short Overview

How do you produce varied language? What kinds of variations exist? When is each kind appropriated?

As shown above, speakers' text variations are based upon the desire to convey different kinds of additional messages to the hearer and to reflect the conversation setting. In order for generator programs to produce similarly varied, appropriate information-bearing text, they must have some means of representing relevant information about the *conversation setting*, and their *interpersonal goals*. These are the *pragmatic* components of language. With respect to language generation, then, the theory is straightforward. Since

different realizations carry different pragmatic effects, pragmatic goals provide a principled way to guide the generation process. Thus, whenever the generator must choose among two or more ways of expressing something, the goals furnish decision criteria. The choice points that enable the topic to be said in various ways encountered during the process of sentence realization include the syntactic issues of language. Some of the more important choice points concern topic collection, organization, and juxtaposition of topics into sentences, selection of aspects of topics to be said, and phrase and word choice.

This PAUTIME works in the following way:




Figure 1. PAUTIME architecture.

In addition to a number of features that represent the conversational setting, PAUTIME uses pragmatic goals to affect the learner in such respects as "partial knowledge," "future behavior," and "new state," interprets new relationships. Initially, the values rep-

representing the conversational setting and the program's *interpreters* and *readers* are set up by the user. In addition, the program is given a topic or topics of conversation, consisting of one or more elements (nodes) in the story representation network. Then the program repeats the following overall cycle. It collects representation elements associated with the topic as possible *elaboration* sentence topics, it orders them, interprets them, and juxtaposes them into *incomplete predicate sentences*; it builds sentences for each topic, choosing appropriate phrases and words to express each element, and determining in what order and which aspects to include (usually as adjectives and adverbial modifiers). After saying the sentence containing the seed for further related topics, and repeats the cycle. However, execution of the cycle can be changed under various settings. For example, under a setting corresponding to *more or less fully planned text*, the program collects all the sentence topics it will say, then finishes all the planning before it produces any text.

1.4 Conclusion

Current Artificial Intelligence work in natural language processing places far too little emphasis on the role of pragmatics. This is a mistake. Not only does this omission limit the power and capabilities of parsers and generators, but it seems to imply that the problem of language will essentially be solved when semantics has been straightened out (assuming, of course, that we already know all the syntax we need). Pragmatics plays a role on many levels of processing — not only in using language and understanding people's goals and *intentions* (relationships, but even *short memory*; for example, in *testing people's retention* of the surface forms of sentences). Keenan, McWhinney & Mayhew '82 describe an experiment that shows that subjects were significantly better at remembering the actual form of sentences that had what they called "high interactional content" (that is, much non-literal, pragmatic information) than remembering sentences with low interactional content. An example of the difference is the following: "I wish you guys would shut up" versus "Please be quiet". Their conclusion is that

findings indicate that the parameters which control the retention of surface form information have yet to be worked out. The present experiment demonstrates the importance of a parameter — that involves the degree to which a statement conveys information about the speaker's intentions, beliefs, and attitudes toward the listener. Our results indicate that the impact of this parameter on memory is more dramatic than that of any of the other parameters which have been examined to date.

Keenan, M., Whinnery, A., Mayhew, S2

The work in this dissertation is based on the recognition that language generation is a goal-driven process, where many of these goals are pragmatic in nature. Nobody has yet attempted to formulate pragmatic goals in terms a generator can plan with and otherwise manipulate, and nobody knows how such goals would interact with a generator's syntactic component. The time is ripe to start examining what kinds of goals are relevant, what strategies achieve them, and what planners and realizers must be like to operate under their control. We must face the fact that, in order to have real, flexible text, we simply cannot do without recourse to the airy world of pragmatics. We need to try out things! It is in this spirit that this work is to be taken.

1.5 The Form of this Dissertation

Given the subject, it is fitting that this dissertation mirrors the generation process. Thus it has the following form: in the first part, discussion centers on pragmatic goals (chapter 2) and on rhetorical goals and strategies for achieving them (chapters 3, 4, and 5). The second part begins with a description of the appropriate nature of the lexicon (chapter 6). Chapter 7 contains a discussion of the kinds of planning required by generation, and illustrates this by describing the program's architecture and the way PAULINE generates a paragraph. Chapter 8, the conclusion, traces the development of AI work in language generation. PAULINE's place in it, and mentions what must come next. Appendix A

contains an annotated trace of a PAULINE session and appendix B contains PAULINE's phrasal grammar.

Unlike most AI dissertations, in which the theory is developed first and the implementation of a program is described later, this dissertation does both in parallel. Each chapter states some problem that arises in generation, develops a pragmatics-based solution, and then describes briefly how the solution is implemented in PAULINE. Thus readers who are not very excited by AI programs may comfortably skip the latter third of each chapter, except with chapter 2.

Chapter 2

Pragmatics

Abstract

Since different realizations of a topic convey different pragmatic effects, the pragmatic aspects of conversations must help to control the choices facing the generator. But, though there has been much discussion about what pragmatics as a field of inquiry is all about, no generally accepted scheme has emerged yet. After a review of some relevant literature, a classification is presented of the kinds of pragmatic information that text can convey. Invariably, however, pragmatic goals are too general to support rules that directly control generator decisions. Thus an intermediate level of goals, specifically attuned to the requirements of language production, are postulated; these goals, called rhetorical goals, determine the slant and the style of text.

2.1 What is Pragmatics? — Some Definitions

There has been much discussion about what pragmatics as a field of inquiry is all about in both the philosophy and linguistics literature (see, say, [Carnap 38, 56, 59], [Morris 38], [Katz 77], [Gazdar 79], [Grice 57], [Gordon & Lakoff 75], [Cole 78], [McCawley 78], [Searle 79], and in summary [Levinson 83]). However, beyond sketching out a number of areas

in which pragmatic considerations play a role, very little agreement has been reached on what exactly pragmatics encompasses and what its relationship is with the other aspects of language and communication in general.

This is not to say that attempts have not been made. In [Levinson 83, p 9], for example, Levinson offers the following definition:

Pragmatics is the study of those relations between language and context that are grammaticalized

Unfortunately, this definition sidesteps the issue of speaker intent. The slurred speech of a drunk or the accented speech of a non-native speaker can also be "grammaticalized", and these aspects also convey non-literal information, but they can only properly be called pragmatic when the speaker uses them to definite purpose.

A definition by Carnap doesn't suffer from this shortcoming ([Carnap 38, p 2]):

If in an investigation explicit reference is made to the speaker, or to put it in more general terms, to the user of the language, then we assign [the investigation] to the field of pragmatics. . . . If we abstract from the user of the language and analyze only the expressions and their designata, we are in the field of semantics. And, finally, if we abstract from the designata also and analyze only the relations between the expressions, we are in (logical) syntax.

In other words,

- syntax = relations between words in a sentence
- semantics = relations between expressions and their designata
- pragmatics = reference to interlocutors and context in communication

Though this provides some idea of what a definition may look like, it lacks a clear enough description of *interlocutor* and *context*: what of "oh!" or "phew!"? These two expressions

illustrate the difficulty of separating semantic content — the meaning of “phew!” — from pragmatic information — the mutual semi-humorous experiencing of a distasteful sensation, with perhaps a plan to avoid it, and so forth. The same problem appears in a definition suggested in [Gazdar 79, p 2]:

Pragmatics has as its topic those aspects of the meaning of utterances which cannot be accounted for by straightforward references to the truth conditions of the sentences uttered. Put crudely:

PRAGMATICS = MEANING – TRUTH CONDITIONS.

What is to be understood by “truth conditions”? Those aspects of semantics that semanticists have succeeded in formalizing? Leaving all the rest to pragmatics? What, for instance, would the truth conditions of sentences such as “can you pass the salt?” include?

The general problem is that no clear distinction exists between pragmatics and semantics. Though attempts to establish the distinction were made (see [Katz 80]) and were countered (in, say, [Jackendoff 81]), the question is not yet resolved. For example, in [Jackendoff 85, p 105]:

... the distinction between “semantic” rules of linguistic inference and “pragmatic” rules of linguistic interaction with general knowledge is less marked than is often supposed. In a theory with an autonomous semantic level, the two kinds of rules involve different levels of mental representation. Here, however, they both are rules for the manipulation of conceptual structures; they deal with the same primitives and principles of combination. If there is a distinction between semantic and pragmatic rules, then, it lies only in the formal manipulations the rules perform on conceptual structure. For example, the principles involved in judging a sentence *true* potentially involve extralinguistic information as well as information within the sentence itself; hence “true” is a pragmatic notion. On the other hand, a judgement that a sentence is *analytic* involves only information conveyed by the sentence itself plus rules of (semantic) inference;

hence “analytic” is a semantic notion. In either case, though, the information conveyed by the sentence is a conceptual structure.

and elsewhere ([*ibid.*, p 106]):

Thus, although a terminological distinction between “semantic” and “pragmatic” notions undoubtedly remains useful, it is an open question whether it is a bifurcation of particular theoretical interest.

This work is based on the belief that it is not of particular interest *in functional terms* (... at least, not until we have a lot more knowledge about semantic representations, speakers’ goals, and speakers’ conceptions of interpersonal relationships). In practice, in any generator written today, no fundamental difference need exist between the nature of processes that make use of so-called pragmatic information and the nature of processes that work with semantic information. Hence no attempt will be made here to provide a definition of pragmatics as a distinct, closed, formalizable field of inquiry. This does not, of course, prohibit identifying a certain body of information and a certain class of considerations as typically pragmatic. This is the sense in which the word “pragmatic” is used here.

2.2 Some Pragmatic Aspects of Conversations

In the philosophy of language, most work on pragmatics focuses on the ways in which text can convey various types of information. The principal areas of interest are:

- *deixis*: references within a discourse to the interlocutors, the time and context, and to the discourse itself. Sentences such as “as long as I’m speaking, you must continue doing that” and “three miles from here, yesterday, a bomb exploded” can only be understood in context.

- *presupposition*: the logical assumptions underlying utterances. Thus after hearing "the king of France is bald" and "flying saucers appeared again last night", the hearer may assume that (as far as the speaker is concerned) there is a king of France and that flying saucers had appeared before.
- *implicature*: the implications (often social, script-based) that can be drawn from utterances, if licensed by certain assumptions. The principal assumptions are called Grice's maxims and cooperative principle (see [Grice 75]). For example, after "where is Mary?", the response "well, her car is in front of the supermarket" is not perverse only if the requester can trust that the responder doesn't in fact know where Mary is, but does know that she travels by car, does know that her car is at the supermarket, that Mary is quite likely to visit the supermarket by car, etc.
- *speech acts*: the effects of stages of utterances (the initial act of production by the speaker, the hearer's understanding of the utterance, the hearer's reaction to it — both his preparation and his actions), and the ways utterances can address certain of these stages and (often only ostensibly) not others. The sentence "I promise to do it" refers both to the speaker's intent to do it as well as to his actual making of the utterance. Very often studied are indirect speech acts, in which the speaker mentions a less important part of a standard sequence of actions (while meaning some other part) in order to give the hearer some leeway, as in the request "can you pass the salt?".

In AI language generation research, most work has been done not on the general *ways* in which pragmatic information can be conveyed, but on *what types* of pragmatic information is used by speakers. In [Cohen 78], Cohen studied the effect of the hearer's knowledge on the selection of appropriate speech act (say, REQUEST vs INFORM OF WANT). Reasoning about the hearer's knowledge in order to plan the inclusion and organization of topics is described in [Appelt 81]; the effect of hearer knowledge on user instruction is described in [Woolf & McDonald 84] and on object description in [Paris & McKeown 87]; the explanation generator of [Swartout 81] had a switch distinguishing between two types of hearer knowledge (either programmer or medical expert). In [Jameson 87], Jameson describes a program that selects

appropriate utterances in evaluative contexts such as job interview situations, based on what effect each utterance is defined to have on the hearer's belief state. [Bienkowski 86] describes automatic elaboration of basic text. Much related work on the structure of discourse uses some pragmatic information, for example, [Grosz & Sidner 85], [Grosz 86], and [McCoy 87].

In addition, a number of general classifications of speaker intent have been made by AI researchers. A very general discussion of speaker goals is given in [Bruce, Collins, Rubin & Gentner 78]. [Johnson & Robertson 81] use goals to model a speaker having a conversation. In [Schank *et al.* 81] the different "points" or speaker intents underlying a statement are analyzed. Other goal classifications can be found in [Carbonell 78] and [De Beaugrande 84]. Furthermore, much sociological and psycholinguistic work has been done in this regard. [Bloomfield 14] mentions the effects of emotional relations on sentences; [Gazdar 80] lists general pragmatic constraints on sentences; [Bühler 34] names some pragmatic aspects of conversations; [Jakobson 60] extends this list. Both [Grimes 75] and [Van Dijk 85] identify a number of pragmatic features and discuss what roles they play in the topic selection, focus, and realization. In the tradition of systemic grammar (see, say, [Halliday 76, 78]), interesting recent work can be found in [Fawcett 80] and [Gregory 82]. The effects of context on utterances is studied in [Clark & Carlson 81], [Clark & Murphy 82]; [Gibbs 79] and [Gibbs 81] discuss the effects of context on the processing of indirect requests. [Osgood 57], [Osgood, Mey & Miron 75], [Osgood & Bock 77] discuss effects of notions such as 'naturalness' and 'vividness'.

What types of pragmatic information can affect language? What additional information can speakers convey in their text? Based on the work mentioned above, the following categorization will be used here (each aspect is briefly discussed below, and some relevant goals are listed):

- interlocutors' factual knowledge
- interlocutors' opinions
- interlocutors' emotional states

- interlocutors' goals (i.e., future behavior)
- interpersonal relationship
- conversational atmosphere (setting)

2.2.1 Interlocutors' Factual Knowledge

The most common speaker goal relates to the hearer's factual knowledge. For example, in

A: "When does the New Haven train leave?"

B1: "The train leaves at 4:15"

each speaker is concerned with facts. (Frequently, though, utterances appear to address factual knowledge but only do so incidentally; in such cases, this goal is subsidiary if present at all; consider

B2: "Oh gosh, the train leaves at 4:15!"

in which the point is a sense of pressure and haste, even though the sentence topic is the same as B1's.)

A speaker can address the hearer's knowledge in one of the following ways:

1. **goal:** access knowledge (i.e., query)
2. **goal:** increase knowledge (i.e., inform, teach)
3. **goal:** reorganize knowledge (i.e., reinterpret, explain, teach)
4. **goal:** "decrease" knowledge (i.e., confuse, obfuscate)

Each of these goals can be achieved in many ways. For example, depending on the nature of the topic, the goal to increase knowledge can take the following forms:

- State fact: "The capital of Switzerland is Berne"
- Describe object: "The vase was blue with a wavy line around the top"
- Relate event: "After that, the victim left the bar..."

Sometimes a topic can be presented in various ways; consider the topic of buttermilk:

- **Instruct:** "Combine:

1 quart 70° to 80° skim milk

$\frac{1}{2}$ cup 70° cultured buttermilk

$\frac{1}{8}$ teaspoon salt

Stir well and cover. Let stand at 70° until clabbered. Stir until smooth. Refrigerate before serving." (from [Rombauer & Becker 75, p 533])

- **Describe:** "BUTTERMILK. Originally this was the residue left after butter making. Today it is usually made from pasteurized skim milk and contains about 8.5% milk solids other than fat. A culture is added to develop flavor and to produce a heavier consistency than that of the skim milk from which it is made." ([ibid.]])

- **Relate:** First, Sue combined some skim milk and some cultured buttermilk. She threw in a dash of salt, stirred it well, and let it stand for a while.

In computer language generation, differences of this kind are achieved by associating with each knowledge goal a set of instructions to indicate which aspects of the topic are appropriate and in what order to say them. Such script-like instructions have been called *schemata* by McKeown (see [McKeown 82]) and Mann ([Mann 84]) and *plans* by Allen and others (say, [Allen & Perrault 79]). The preponderance of knowledge-related goals in conversations has sparked much work on recognizing them ([Allen & Perrault 79], [Litman & Allen 84]), expressing them by means of speech acts ([Cohen & Perrault 79], [Cohen 78]), describing their elements ([McKeown 82], [Mann 84], [Rösner 87]), constructing plans by reasoning about the world and beliefs ([Appelt 81]), planning while teaching ([Woolf & McDonald 84]), and extemporizing ([Bienkowski 86]). In [Rich 79] Rich describes a librarian-modeling system that builds a stereotype of the user and uses information regarding the user's likely knowledge level and interests to select what to say. PAULINE's topic collection plans are described in chapter 4 and its planning methods in chapter 7.

2.2.2 Interlocutors' Opinions

Almost all spoken and written communication contains the interlocutors' opinions, whether explicit or implicit. In our daily lives, we constantly face a barrage of persuasive text from blatant advertising to subtle cozening. In fact, our biases sneak into what we say so easily and so often that producing genuinely unslanted text can be quite a problem! Certainly (a) seems a more natural description of a fight than (b) does:

(a) After Jim, without any provocation, bumped into Mike, Mike just gave him a gentle tap back. Then Jim hit Mike as hard as he could. Poor Mike just had to knock Jim out in self-defense.

(b) After Jim bumped into Mike, Mike hit him back once. Then Jim hit Mike moderately hard, and Mike knocked Jim out.

In order to make appropriate use of the stress words, adjectives, and all the other possibilities for injecting affect into its text, a generator has to have explicit goals and strategies that prescribe when and how to slant its text. For a given topic about which the hearer is known to have some opinion, such goals are:

5. **goal:** make the topic seem good or bad, contrary to the hearer's opinion
6. **goal:** make the topic seem good or bad, bolstering the hearer's opinion

Pursuing affect goals can have side-effects that serve or hinder other goals. For example, when the speaker's desired opinion opposes that of the hearer, the latter may become angry, and eventually dislike the speaker. To forestall this, the speaker must attempt argument or persuasion. Analysis of the relationships among elements of arguments was done in the context of understanding arguments by [Birnbaum 85], [Birnbaum, Flowers & McGuire 80], and a study of manipulating them during arbitration was done by [Sycara-Cyranski 85]. In [Jameson 87], Jameson describes a program that selects appropriate utterances in evaluative contexts such as job interview situations. Strategies to achieve the goal of convince the hearer to change his opinion are discussed in chapter 4.

2.2.3 Interlocutors' Emotional States

Many utterances spring from the desire to create a given emotion in the hearer. Though any discussion of emotions is beyond the scope of this work, one emotional state — that of anger — can be achieved relatively easily by a generator, by the simple method of inverting certain of its rules of status and respect, as described in chapter 5. Discussions of politeness appear in [Goody 78], [R. Lakoff 77], and [Clark & Schunk 80].

2.2.4 Interlocutors' Goals: Altering Future Actions

Many language interactions occur when a speaker wants to affect the hearer's future actions; for example,

Keep off the grass!

Often, such desires arises in service of the speaker's goals for his own future. For example, if the speaker wants to know the time, he must activate in the hearer the goal to inform him (the speaker); this causes requests such as

"Excuse me, do you have the time?"

All requests, questions, orders, and warnings attempt somehow to affect the hearer's goals, plans, and actions. They are, in short, the applicable elements of the planbox of the goal D-CONT in [Schank & Abelson 77]:

- **Request:** "Where is the station?"
- **Make indirect request:** "You are standing on my foot"
- **Order:** "Go wash the floor!"
- **Bargain:** "If you pay me \$25 I'll cross the street on my hands"
- **Threaten:** "If you don't give me some ice-cream I'll rub sand in your hair"

as well as

- **Warn:** "Beware of the quicksand!"

With respect to the hearer's goals, a generator should have at least the following goal:

7. **goal:** activate or deactivate a specific goal in the hearer

Much philosophical work has been done on the nature and interpretation of requests (see especially [Searle 69, 75] on indirect speech acts, and [Levinson 83] for a summary). With respect to computer systems, Cohen describes the selection of appropriate speech act (REQUEST vs INFORM OF WANT) in [Cohen 78]; some criteria that determine the selection of the form of a request are studied in [Kempen 77] and [Herrmann & Laucht 78]; a discussion of the relation between the degree of directness of a request and the speaker's certainty of its being granted appears in [Laucht & Herrmann 78].

2.2.5 Interpersonal Relationships

A large amount of psycholinguistic and sociological work studies the way language expresses relationships between interlocutors. In [Wish, Deutsch & Kaplan 76], Wish, Deutsch and Kaplan scaled subjects' responses and found that four dimensions captured most of the distinctions subjects considered important:

- regard: positive negative
- relative social status: dominant inferior
- intensity: intense superficial
- formality: formal informal

A similar study by Joncas ([Joncas 72]) yielded three distinct dimensions: regard, relative social status, and formality. In [Schank & Abelson 77], Schank and Abelson discuss life

themes (the long-term goals people typically have). Of these, interpersonal themes give rise to goals of various levels of specificity that can affect language generation. For example, the speaker's goal CAUSE Y BE MENTAL-STATE (HAPPY) can activate an appropriate generation goal; the goal A-RESPECT (Y) can activate the generator goal to increase Y's social status by using appropriate address forms and expressions. Straker, in [Straker 80], tested the effects of intimacy, setting, and social status on language dialect and style with 28 students speaking both standard American and Black English. By varying the following three parameters

- Interlocutor status: dominant - equal (instructor - student)
- Topic formality: formal - informal (subject's aspirations - worst experience)
- Setting formality: formal - informal (office - lounge)

she found that the interlocutor and topic significantly determined language style: the subjects used Black English only in intimate conversations when the interlocutor was a student too. [Schank *et al.* 81] analyze the various kinds of "points" a speaker can make with an utterance in a conversation, and [Johnson & Robertson 81] use some of these in a computer model. [Brown & Levinson 78] describe aspects of language use that depend on social relationships.

This discussion will address three ways in which speakers and hearers relate: interpersonal emotion, interpersonal distance, and relative social status.

(a) Hearer's Emotions toward the Speaker

Usually, speakers have the goal, at some low level of priority, of making their hearers like them. This goal gives rise to a number of goals that most speakers have active during their conversations and that act as default goals when more specific ones are lacking. Sometimes, however, the goals to affect the hearer's emotions toward the speaker are of primary importance, and determine topic choice and especially topic presentation; one typical result is,

of course, flattery. Again, though this is not the place to discuss emotion, some emotional effects can be achieved by relatively straightforward generation techniques. Two typical goals are:

- 8. **goal:** make the hearer respect the speaker
- 9. **goal:** make the hearer like or dislike the speaker

(b) Distance between the Hearer and the Speaker

Similar to but distinct from the hearer's feeling toward the speaker is his perception of the distance between them. The speaker can have the goal to have the hearer like him, yet not feel very close to him; this is a goal, for example, that many people in positions of authority have with respect to their subordinates. This goal may cause a speaker to be formal yet polite, as is discussed in chapter 5. The generator's pragmatic goal of distance should be:

- 10. **goal:** make the hearer feel closer to or more distant from the speaker

(c) Relative Status of the Hearer to the Speaker

The relative social status of interlocutors has a large effect on their text. For example, after the following interaction, what effects do A's responses 1, 2, and 3 have on B?

A: "How are things?"

B: "Not good"

A: "How come?"

B: "I'd rather not discuss it"

(1) A: "Hey just watch who you are talking to!"

(2) A: "Oh, sorry: please excuse me"

(3) A: "Ok"

Clearly, in (1) A wants to establish his dominance over B, whereas in (2) his apology indicates his subordination to B, and in (3) he signals neither. Studies on the ways in which

languages communicate status information abound; for example, see [Kuno 73], [Harada 76], [Gasser & Dyer 86] on Japanese deictic honorifics. Experiments described in [Straker 80] showed the effect of social differences on speakers' choice of dialect. Typically, a speaker's status-related pragmatic goal affects the way the topic is expressed rather than which aspects of the topic are chosen to be included. In a generator, the status goal can take the following form:

11. **goal:** make the hearer feel socially inferior to, equal with, or dominant over the speaker

2.2.6 Conversational Atmosphere

Each conversation takes place in a setting which, by virtue of factors such as privacy, topics of conversation typically addressed or socially prescribed, and amount of background disturbance, has a tone or level of formality and intimacy. The form of the language used depends as much on the setting as on the interlocutors' desires: it is much harder to be formal at a picnic or at a football game than in an office or at a conference; it is much harder to be uninhibited and intimate in a subway car than at home. Still, the speaker can manipulate to some extent the atmosphere of the conversation by selecting appropriate topics and by employing appropriate phrasing; consider, for example, the difference between

(a) "May I introduce Maria...she is Pete's girlfriend."

"Pleased to meet you! Is there any special reason you are here in Sydney?"

(b) "This is Pete's girlfriend Maria."

"Oh hi! So what brings you to Sydney?"

In this category a generator should have a pragmatic goal of the following form:

12. **goal:** make the tone formal, informal, or intimate

Much work by sociologists, anthropologists, and psycholinguists describes the characteristics of various settings and the appropriate language styles; for example, [Irvine 79] and

[Atkinson 82] describe formal events in various cultures; [Brown & Levinson 78] discuss the use of honorifics in formal situations; [Levinson 83] presents an overview of some relevant literature.

2.3 PAULINE's Pragmatic Settings and Goals

The above discussion is very vague. Each individual aspect mentioned, and its relation to the generation process, is a large and complex field of study; certainly nobody is in a position today to formulate the pragmatic aspects of language concretely enough to model in a computer program. Nonetheless, if we want to generate flexible and appropriate language by computer, we have to take into account pragmatic considerations in some form or other. As a first approximation to characterizing the pragmatics of conversations, PAULINE was simply given a list of features that are similar to those discussed above. The (pragmatic!) justification of these features is that they are the kinds of features necessary to make a generator produce these types of text. No additional claims are made about the completeness or adequacy of this categorization.

In the representation of pragmatics used here, each feature is given a fixed number of distinct values. Usually, the values lie on some intuitive scale; in a few cases, two or more scales are conflated and the result is merely given as a set of distinct values; this can eventually be refined. PAULINE's characterization of its **conversation setting** is (where default values are the middle settings):

- Conversational Atmosphere (setting):
 - **time** -- *much, some, little*
 - **tone** *formal, informal, festive*
 - **conditions** *good, normal, noisy*
- Speaker:

- knowledge of the topic — *expert, student, novice*
- interest in the topic — *high, normal, low*
- opinions of the topic — *good, neutral, bad*
- emotional state — *happy, angry, calm*

- Hearer:

- knowledge of the topic — *expert, student, novice*
- interest in the topic — *high, normal, low*
- opinions of the topic — *good, neutral, bad*
- language ability — *high, normal, low*
- emotional state — *happy, angry, calm*

- Speaker-Hearer Relationship:

- depth of acquaintance — *friends, acquaintances, strangers*
- relative social status — *dominant, equal, subordinate*
- emotion — *like, neutral, dislike*

In addition, PAULINE has been given the following interpersonal goals:

- Hearer:

- affect his knowledge — *teach, inform, confuse*
- affect his opinions of topic — *switch, no effect, reinforce*
- involve him in the conversation — *involve, no effect, repel*
- affect his emotional state — *anger, no effect, calm down*
- affect his goals — *activate, no effect, deactivate*

- Speaker-Hearer Relationship:

- affect hearer's emotion toward speaker — *make respect, like, dislike*
- affect relative status — *make feel dominant, equal, subordinate*
- affect interpersonal distance — *make intimate, close, distant*

2.3.1 From Pragmatics to the Generation Process

Rhetorical Goals and Generator Decisions

Though the pragmatic aspects of the conversation help determine the speaker's text, most do not do so directly, since they are too general to be attuned to the requirements of language production. As a result, attempts to write down rules that relate pragmatic aspects to generator production decisions are doomed to failure; inevitably, such attempts quickly become bogged down in minutiae and produce rules with little credibility. For example, what is the effect on sentence length if the speaker is socially dominant over the hearer? If the speaker is antagonistic toward the hearer, should he make active or passive sentences? Is it right to say that if the speaker wants to impress the hearer he should always select formal words and phrases?

Yet, of course, pragmatic aspects do influence text production. Therefore, rules must exist that relate these aspects to the generation process. These rules must depend on the pragmatic aspects and must interact with the generation process in order to produce text that serves the speaker's goals. Since the interpersonal goals are too far removed from the syntactic concerns of language to provide such rules, *there must exist a number of other goals expressly designed for this purpose.*

These goals will be called *rhetorical goals*. They act as intermediaries between, on the one hand, the speaker's interpersonal goals and other pragmatic aspects of the conversation, and, on the other, the syntactic decisions a text producer has to make. These entities are called here *goals* rather than *strategies* specifically to emphasize their independence from the system's interpersonal goals. After all, any identifiably distinct collection of information that is activated to guide the behavior of the system toward a desired specific final state can be called a goal; in practice, we dignify those collections that we consider somehow important or "natural" starting points by calling them goals, and the rest we simply call strategies or plans. In PAULINE, rhetorical strategies are only applied in decisions if the rhetorical goal controlling them has been activated.

The production of a sentence involves a large number of decisions (in the form of selecting from a set of alternatives); a typical sentence, such as this one, can require about 50 decisions (pretend you are a generator and count the the number of ways you can say this!). Consider, for example, the following simple representation elements (ATRANS represents the transfer of control over the OBJECT; this is part of Conceptual Dependency Theory, a system for representing semantic information, developed in [Schank 72, 75, 82] and [Schank & Abelson 77]):

ATRANS	ATRANS
ACTOR John	ACTOR Mary
OBJECT book	OBJECT money
TO Mary	TO John
FROM John	FROM Mary
LOCATION store	LOCATION store
TIME yesterday at T1	TIME yesterday at T1

Any reasonably powerful generator must make at least the following six types of decision:

- **topic choice:** collect aspects of topic and related aspects as sentence topics
- **topic organization:** find appropriate groupings and interpretations of candidate topics; find appropriate ways to juxtapose candidates in multi-predicate phrases; find ways of expressing relations among topics. For example, the two representation elements can be related in various ways: by simple conjunction (using "and", as in "John gave Mary the book and she gave him the money"); by merging them (using "sell", as in "John sold Mary the book"); or by subordinating one element to the other (using a relative clause, as in "Mary got the book from John, to whom she gave the money")
- **sentence inclusion:** select appropriate aspects of each sentence topic. In the example, the inclusion of John, Mary, and the book, as well as of the time and location,

must be determined. Some variations result in "John sold the book yesterday", "Mary bought the book in the store", etc.

- **sentence organization:** select appropriate subject, pre-sentence adverbial clauses, verb, predicate clauses, etc., and order them. For example, both John and Mary can be selected as sentence subject; this decision will help determine the verb ("John sold Mary the book" and "Mary bought the book from John"). In addition, the position and order of the adverbial clauses can be "yesterday, John sold Mary the book in the store" or "John sold Mary the book in the store yesterday", etc.
- **clause/noun group inclusion and organization:** determine which aspects of each clause topic to include and determine their order. Different inclusion decisions produce noun groups such as, for instance, "the big blue book", "the book", and different ordering decisions "big blue book" and "blue big book"
- **word choice:** select appropriate words and phrases. For example, "tome" or "novel" can be used instead of "book"; "boutique" or "shop" for "store"

The simplest existing generators, of course, perform these decisions by having only one available option. However, as soon as the generator is given the ability to realize the topic in more than one way, it has to be able to make its choice in a principled way. Since different realizations convey different pragmatic effects, the pragmatic aspects of conversations *must* help determine the choices. As the agents of these aspects, rhetorical goals supply the criteria by which these decisions are made; if the final text communicates any additional information at all, it is due to the control exercised over the generator by the set of active rhetorical goals. This is the relationship between the rhetorical goals and the text production system.

Pragmatics and Rhetorical Goals: The Case for Their Independent Existence

On the other hand, the relationship between the rhetorical goals and the pragmatic aspects of the conversation is not so clear-cut. Pragmatic-based language generation would be

simple if each rhetorical goal reflected one and only one interpersonal goal or conversational aspect. In this case, each rhetorical goal would simply be a repository for the generator-specific knowledge required to express its pragmatic partner. But the pragmatic aspects of conversations are not independent; they influence each other. This fact makes the rhetorical goals more complex. To see why, note that a single rhetorical goal can express opposite pragmatic aspects under different conditions. For example, if the speaker has the goal to make the hearer feel close to him, he may activate the rhetorical goal to be humorous (say, by choosing funny words and by selecting funny topics). Usually this will work well, but it will backfire if the hearer has just heard of his (the hearer's) mother's death. In this case, an appropriate rhetorical goal is the goal to be serious and slightly formal — which, under normal circumstances, would tend to alienate him. Different rhetorical solutions can achieve the same pragmatic goal under different circumstances.

As it is, a pragmatic effect is seldom the result of a single rhetorical goal; combinations of rhetorical goals act in concert to produce pragmatic results. For example, low **formality**, high **force**, and high **partiality together** have an effect on the text that is distinctly pragmatic and can be characterized as *no-nonsense*. (Similarly, *blather* is the result of high **formality**, low **force**, and low **partiality**...) Therefore, rhetorical goals cannot simply be paired one-to-one with pragmatic aspects, unless a distinct goal is defined for each possible combination of aspects. Rhetorical goals have to be independent goals in their own right.

A second reason for defining rhetorical goals as independent carriers of pragmatic information is more practical. At this time, the field of pragmatics contains a number of very complex unsolved issues. In order to generate language with pragmatic effects today, we can use rhetorical goals as a starting point. Whatever form the eventual pragmatic solutions take, they will be able to interface with a pragmatics-sensitive generator through a set of rhetorical goals. These intermediate goals then are a reasonable level to which both generator builders and researchers in pragmatic issues can relate in order to find common ground.

Furthermore, the notion of rhetorical goals as independent entities is useful from an AI/programming/engineering point of view: they provide a useful level of organization for

certain types of criterial information. In practise, the generator builder sometimes finds that when a number of generator decisions vary together, the text has an unexpected pragmatic import. He can then assemble the relevant information and define a new rhetorical goal to take care of the issue. For example, consider the honesty of a generator that says "tap" when its input representation is ACTION HIT with aspect DEGREE HARD. If the generator is to slant its text in order to support the hitter, and its verb choice strategies prescribe the use of "tap" — is it lying or not? Must the generator stick to the 'letter of the representation'? Furthermore, what should it do about the use of adjectives and adverbs: may it say "tap lightly"? And what about sentence topics as a whole: may it suppress topics that hinder its goals? In some conversations, the generator must be scrupulously honest; in others, it may have more leeway. The generator designer can then group together the relevant decision strategies and activation criteria and define a new rhetorical goal called **honesty**.

The advantages of identifying and using a set of intermediate goals should be obvious. Not only do they seem intuitively plausible, but they furnish a place to make explicit, collect, and organize many generator decisions and design characteristics that most generators have left implicit or avoided altogether.

Each rhetorical goal causes characteristic effects in the text. Different combinations of rhetorical goals result in differences in textual content and form. Some rhetorical goals are achieved by slanting the text; others find their expression as the style of the language. Through slant and style, the speaker can communicate additional information that the hearer can interpret and respond to. The rhetorical goals of opinion are described in the next section and discussed in chapters 3 and 4; following them, some rhetorical goals of style are listed. They are discussed in more detail in chapter 5.

2.4 Rhetorical Goals of Opinion

A very common speaker goal is to alter the hearer's opinion about a topic. This is the goal, for example, in all advertising, one of the primary goals in biased reporting, and one of the goals in soliciting money or help. Usually, at least two opinions exist — for the topic

(sympathetic to it) and against (antipathetic to) it. Often two parties are involved — “our side” and “the opponent”. Language users have a number of techniques for manipulating the presentation of the topic in order to slant it one way or the other. Consider again the slanted texts quoted in the previous chapter:

It is time, in the best interests of all concerned, to settle the strike. It is our understanding that the University administration, as well as the Union, has received a document entitled ‘A Statement of Purpose by the Coalition to End the Strike’. We appreciate the spirit of the document. Clearly, the community earnestly desires and needs a settlement, so that Yale can get back to what it is supposed to be. Our members earnestly want a settlement. You have said that you do, too, and we are prepared to take you at your word. We are willing to compromise significantly to achieve a settlement. Therefore, we propose the immediate resumption of negotiations on a daily basis.

(From: open letter from Yale university’s clerical and technical workers’ labor union negotiating committee, November 9, 1984)

I write with great disappointment following Local 34’s action in calling a strike against the University. The University negotiating team has made concerted efforts, lately with the help of the Mediator, Eva Robins, to find common ground and to bring about a fair and reasonable settlement of the outstanding issues in a manner satisfactory to both parties. But the agreement has not been achieved.

(From: open letter from Yale president Giamatti to the university community, September 26, 1984)

Some slanting techniques the speakers used caused them to say the following:

- “in the best interests of all concerned” — they care about the university’s goals
- the union “earnestly desires and needs a settlement” — they want to settle

- they "are willing to compromise significantly" they want to settle
- they will "take [Yale] at [its] word" they are trusting
- they "propose an immediate resumption of negotiations" they want to negotiate
- Giamatti writes "with great disappointment" he is reasonable; the union isn't
- his team has "made concerted efforts" they tried to negotiate
- "in a manner satisfactory to both parties" he cares about the union's goals
- the agreement "has not been achieved" the union won't settle

On April 4, concerned Yale students constructed a shantytown on Boesak Plaza as a reminder to those in Woodbridge Hall (and all over campus and the community) that Yale is complicit (sic) with the system of apartheid that creates shantytowns where thousands of blacks are forced to live in squalor and fear. The shantytown, Winnie Mandela City, served as a focal point of education concerning South Africa and Yale's investments there. At 5:30 am on April 14 the Yale Administration had the shantytown torn down and had 76 students and community members who were defending the shanties arrested. After a huge outcry, the Administration allowed the shanties to be rebuilt. We will not be silenced; we will continue to challenge the University on their moral failure. (From: protester literature; the protesters renamed the plaza after the South African churchman Allan Boesak)

On April 4, a small group of students took over Beinecke Plaza and built some shanties; they wanted to force Yale to sell its stocks in companies with branches located in South Africa. The university asked the students to move the shanties to another location, but the students refused. The university then granted them permission to occupy the plaza until the end of the week, so that they could be there to be seen by the university's trustees, the Yale Corporation.

at their meeting. But even after the meeting, the students refused to leave the plaza, and police had to clear the shanties. Later, the university relented, and gave them permission to rebuild the shanties. It also announced that it would send a fact-finding mission to South Africa.

(Speaker: anti-divestment Yale student)

Some more slanting techniques the speakers used:

- “concerned Yale students” — they care about others’ prosperity
- “constructed a shantytown... as a reminder to Yale” — they are not aggressive
- Yale “had 76 students... arrested” — Yale is aggressive and nasty
- “a huge outcry” — they have much popular support
- “a small group of students” — they have little support
- “took over Beinecke Plaza” — they overstep the bounds of propriety
- “they wanted to force Yale...” — they are aggressive, coercive
- “police had to clear the shanties” — the police were forced into action

The main slanting goal activates subgoals, the *rhetorical goals of opinion*, which are served by the slanting techniques. Techniques for making one party (“our side”) look good can be classified as follows (they are described in more detail in chapter 7):

- **Our goals:** Show how our side has good goals, by describing how (a) we help other people; (b) we want a solution to the conflict; and (c) our goals are good according to accepted standards.
- **Our actions:** Explain how our side does good actions to achieve the goals: (a) the actions are not unreasonable or nasty; (b) they are good according to accepted

standards; and (c) they are performed in the open. In addition, describe (d) our side's response to the opponent: which negotiations that have taken place and how we have moderated our demands

- **Our claim:** State outright that our side is good
- **Our reactions:** Show our reasonable reaction to their actions, such as that we were (a) disappointed; (b) hurt; or (c) outraged; or else (d) satirize their actions
- **Our support:** Show how other people believe that we are good, by describing (a) their active support and (b) their statements and recommendations to that effect

The inverse goal to show how bad the opponent's side is can be similarly subclassified:

- **Their goals:** Show how their side has bad goals, by describing how (a) they are only in it for their own benefit; (b) they don't really want a solution to the conflict; (c) their demands are beyond reasonable expectations; and (d) their goals are immoral *and unfair according to accepted standards*
- **Their actions:** Explain how their side does bad actions to achieve their goals: (a) they started the whole affair; (b) their actions are ugly, distasteful and overstep the bounds of propriety; (c) the actions are aggressive and inciting; (d) they coerce other people into doing things for them; (e) they disseminate false or misleading information; and (f) they have a hidden agenda. In addition, describe their response to our overtures: (g) they won't negotiate; and (h) they won't moderate their demands
- **Our claim:** State outright that their side is bad
- **Their reactions:** Show their unreasonable reactions by saying that they are (a) nasty and spiteful; (b) gleeful at our misfortune and suffering; and (c) intransigent and unconciliatory
- **Their support:** Show how nobody likes them (a) their events are not well-attended; (b) people attack them publicly; and (c) they claim to have more support than they really have

The subgoals in both classifications suggest specific sentence topics under appropriate circumstances. Clearly, not all these subgoals are appropriate in every conversation about which opinion differs! Thus each subgoal must be associated with conditions for its activation; and therefore they can be thought of as inferences. The inference process by which given topics are interpreted as other concepts is described in chapter 3, and the process by which additional topics are introduced is described in chapter 7.

2.5 Rhetorical Goals of Style

In addition to opinions, text can convey a lot of other information. Consider the following example: When, in [Wodehouse 79, p 37], the butler Jeeves says to his master Wooster

The scheme I would suggest cannot fail of success, but it has what may seem to you a drawback, sir, in that it requires a certain financial outlay.

and Wooster paraphrases this to a friend as

He means... that he has got a pippin of an idea, but it's going to cost a bit.

we understand that the former is urbane, formal, and perhaps a little smug, while the latter is young and trendy. By varying the *style* — by making Jeeves's text highfalutin and Wooster's slangy — the author has communicated far more to us than the literal content of the forty-six words.

In order to produce pragmatic-based, goal-directed language, then, we have to understand style: what it is, what effects various styles have on the hearer, and what information various styles convey.

Classifying all possible styles of text is an impossible task. Since one can imagine text characteristics that fit almost any adjective (for example: *heated* text: short, explosive sentences, full of opinions, forceful language; *greasy* text: devious, subtle presentation, dishonest flattery). And clearly, with any well-defined set of stylistic primitives, such circular

definitions are worthless). In order that generators be able to produce pragmatically varied text, we require a theory of style that provides components that can be used in programs and from which we can build various styles (even heated and greasy ones).

A study of some of the major handbooks of good writing (such as [Weathers & Winchester 78], [Birk & Birk 65], [Payne 69], [Hill 1892], [Loomis, Hull & Robinson 36], [Baker 66], [Cowan & McPherson 77], [Strunk & White 59], [Willis 69]) indicates that the authorities agree on a few such common broad-based features in their discussions of style. For example, some of the more complete categorizations of various features of style are:

- *formality, texture, emphasis* (Weathers & Winchester)
- *coherence, concreteness, economy, emphasis, formality, tone, unity, variety* (Birk & Birk)
- *clearness, force, ease, unity* (Hill)

These features they describe in terms of characteristics of complete paragraphs of text. Unfortunately, this descriptive approach is of very little use in a theory of language production, since it never makes clear why and how each style is formed out of words; nor does it indicate any systematicity behind the classification of the stylistic features.

In contrast to such descriptions, a functional approach is to describe styles in terms of the decisions a generator has to make. The decision-based approach enables a more concrete description of each style and its relation to other styles.

Just as the rhetorical goals of opinion determine the slant of text by controlling generator decisions, the *rhetorical goals of style* determine the style. These goals control such traditional notions as formality, force, and respect. Having been discovered during the construction of PAULINE (rather than through abstract reasoning or psycholinguistic experimentation), these goals are motivated empirically—when you vary the decisions you make during generation, certain types of decision group together and form stylistically coherent text, and other types, when grouped, produce text that is incoherent (11). The

coherent groupings conform to traditional stylistic concepts. The classification of stylistic goals presented here is not the only possible one; many groupings are open to reorganization and reformulation. It is not complete or completely consistent. The claims made here are about: *function of style* — the expression of rhetorical goals in order to achieve pragmatic goals in the text; and *method of definition of style* — defined as constraints on the decisions the generator has to make.

In this dissertation, all PAULINE's rhetorical goals will be prefixed by **RG:**. Thus, for example, **RG:formality** refers to the collection of strategies that control the generation of formal or informal text. PAULINE's rhetorical goals of style are contained in the following list (names are somewhat whimsical; this is to suggest their function without identifying them too closely with traditional stylistic terms). This list does not contain all possible rhetorical goals, since such a list is impossible to make: every speaker has an idiosyncratic set of goals and techniques for manipulating language. However, this list contains of the common rhetorical styles; most other rhetorical goals are refinements and extensions of them. Anybody is welcome to define his own particular *heated* and *greasy* text styles in this manner, either in terms of generator decisions or in terms of the styles described here. The goals are:

- **RG:formality** (*highfalutin, normal, colloquial*): Highfalutin language is used for speeches and toasts
- **RG:simplicity** (*simple, normal, complex*): Simple text has short sentences and easy words
- **RG:timidity** (*timid, normal, reckless*): Willingness to include opinions at all
- **RG:partiality** (*impartial, implicit, explicit*): How explicitly opinions are stated
- **RG:detail** (*details only, interpretations, both*): Too many details can be boring to non-experts
- **RG:haste** (*pressured, unplanned, somewhat planned, planned*): When there's little time, you speak fast

- **RG:force** (*forceful, normal, quiet*): Forceful text is energetic and driving
- **RG:floridity** (*dry, neutral, flowery*): Flowery text contains unusual words
- **RG:color** (*facts only, with color*): Colorful text includes examples and idioms
- **RG:personal reference** (*much, normal, none* — two ranges, for speaker and hearer): Amount of direct reference to the interlocutors
- **RG:openmindedness** (*narrow-minded, openminded*): Willingness to consider new topics
- **RG:respect** (*arrogant, respectful, neutral, cajoling*): Techniques for communicating relative status

Each of the rhetorical goals mentioned above is implemented in PAULINE. Each goal is activated by criteria that depend on the program's initial set of pragmatic values and goals; in turn, each goal activates a number of strategies that guide the generator during the planning and realization of text. This guidance takes the form of suggestions at choice points, whenever the generator encounters more than one topic-related, phrasal, or syntactic option. The rhetorical goals **RG:detail** and **RG:color** are discussed in chapter 3; **RG:partiality** and **RG:timidity** appear in chapter 4; and **RG:formality**, **RG:haste**, and **RG:force** are discussed in chapter 5.

2.6 Conclusion

In order to begin to study how pragmatics is used in generation, a number of rather crude assumptions must be made about plausible types of goals of speakers and about the relevant characteristics of hearers and of conversational settings. The specific pragmatic features used by PAULINE are but a first step. They are the types of factors that play a role in conversation; no claims are made about their literal veracity. Similarly, the strategies PAULINE uses to link its pragmatic features to the actual generator decisions, being dependent on the definitions of the features, are equally primitive; again, no strong claims are

made about their existence in people in exactly the form shown. However, in even such a simple theory as this, certain constraints emerge, and these constraints, I believe, hold true no matter how sophisticated the eventual theory is. The constraints pertain primarily to the organization of pragmatic information in a generator: (a) the fact that pragmatic and interpersonal information is too general to be of immediate use; (b) the resulting fact that intermediate strategies, here called rhetorical strategies, are required to run a generator; (c) the fact that, as described in chapter 7, in a model of generation that incorporates these goals rhetorical planning and realization are interleaved processes, where the interleaving takes place at the choice points (this view supports the standard top-down planning-to-realization approach, as well as a bottom-up approach, in which partially realized syntactic options present themselves as opportunities to the rhetorical criteria, at which point further planning can occur). This design can be called a limited-commitment planner that satisfies its pragmatic goals opportunistically.

Chapter 3

Interpretation

Abstract

The computer maxim *garbage in, garbage out* is especially true of generation. When a generator slavishly follows its input topics it usually produces bad text. One remedy is to give the generator the ability to decide what topics to include and at what level of specificity — that is, the ability to interpret its input as instances of other representation elements. Since interpretation requires some inference, generators must be able to exercise some control over the inference process. Some general strategies of control and some specific techniques, geared toward achieving pragmatic goals, are described in this chapter.

3.1 The Problem

Simply put, the generator's task, for a given sentence topic, is to find a form of expression — either a syntactic rule or a phrase — that will enable it to select and to order aspects of the topic in order to build a sentence. The straightforward approach is to define a fixed correspondence between topic representation types on the one hand and grammatical rules and lexical elements on the other. This approach has a flaw: the results are invariably

bad or boring. How bad, of course, depends on the representation, but anything detailed enough to be useful for other purposes, such as learning or diagnosing, simply does not make great prose in practise. A good example is furnished by the JUDGE texts described in chapter 1. In this example, the generator's input consists of a list of topics, where each topic describes some episode in a fight between two people. Straightforward generation (also done by PAULINE) produces:

(a) FIRST, JIM BUMPED MIKE ONCE, HURTING HIM. THEN MIKE HIT JIM, HURTING HIM. THEN JIM HIT MIKE ONCE, KNOCKING HIM DOWN. THEN MIKE HIT JIM SEVERAL TIMES, KNOCKING HIM DOWN. THEN JIM SLAPPED MIKE SEVERAL TIMES, HURTING HIM. THEN MIKE STABBED JIM. AS A RESULT, JIM DIED.

This example is an extreme case because it contains only two main representation types, ACTION and STATE, which can relate in only one way, RESULT. When the generator knows only one way to express this combination, what more can we hope for? Though the problem is less apparent in stories that contain more representation types (and hence a larger variety of sentence patterns), it still is a problem.

Correcting this inflexibility seems straightforward. Though there is nothing wrong with the sentence form used above, namely

[[SAY-TIME #TIME] [SAY-SENTENCE #ACTION] , [SAY-PARTICIPLE #STATE]]
one can add to the grammar a few more sentence forms expressing actions and their resulting states, as well as some more time words and verbs, and then make the generator cycle through its options whenever it encounters a choice point:

(b) 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, HURTING HIM. AFTER THAT, MIKE STABBED JIM. AS A RESULT, JIM DIED.

Yet this produces no real improvement! Clearly, simply extending the number of phrase patterns for each representation type does not solve the problem. When we speak, we do a lot more than simply cast input topics in various forms: of the fight, for example, a person might say the following:

(c) JIM DIED IN A FIGHT WITH MIKE.

(d) AFTER JIM BUMPED MIKE ONCE, THEY FOUGHT, AND EVENTUALLY MIKE KILLED JIM.

(e) AFTER JIM BUMPED MIKE ONCE, THEY FOUGHT, AND EVENTUALLY HE WAS KNOCKED TO THE GROUND BY MIKE. HE SLAPPED MIKE A FEW TIMES. THEN MIKE STABBED JIM, AND JIM DIED.

Illustrated this way, the problem seems rather simple. Obviously, the solution is to group together similar enough topics, where the similarity criterion can be varied depending on external factors, and then to generate the groupings instead of the individual actions. Doing this, PAULINE produced variants (c), (d), and (e) by grouping together contiguous actions of similar force. (In the first variant, all actions were grouped together; in the second, all actions more violent than bumping but less violent than killing; and in the third, the grouping resulted from defining four levels of violence : bumping, hitting and slapping, knocking to the ground, and killing.)

How, then, do we group together input topics? What are appropriate grouping criteria? Clearly, though it improves the JUDGE examples, the technique of grouping actions by levels of force is very specific and not very useful. However, when "group" is used in a wider sense to mean "interpret", this technique becomes both difficult and interesting, and provides a very powerful way to increase the expressive flexibility and text quality of a generator. So the questions are: what interpretation/grouping criteria are general and still useful? When and how should the generator interpret input topics? How should it find appropriate grouping criteria?

3.2 Interpretation in Generation

Consider again the example of the fictitious primary between Carter and Kennedy from chapter 1. In straightforward generation of the outcome for each candidate, PAULINE says:

(f) IN THE PRIMARY ON 20 FEBRUARY, CARTER GOT 20515 VOTES. KENNEDY GOT 21850.

However, since PAULINE can notice that the two outcomes relate to the same primary, it can say either of the following instead:

(g) IN THE PRIMARY ON 20 FEBRUARY, KENNEDY BEAT CARTER BY 1335 VOTES.

(h) IN THE PRIMARY ON 20 FEBRUARY, CARTER LOST TO KENNEDY BY 1335 VOTES.

(or any of a number of similar sentences with "beat", "win", and "lose"). But why stop there? If PAULINE examines the input further, it can notice that Carter's current delegate count is greater than Kennedy's, that this was also the case before the primary, and that the primary is part of a series that culminates in another election, the nomination. In other words, PAULINE recognizes that what happened in this primary was that

(i) IN THE PRIMARY ON 20 FEBRUARY, KENNEDY NARROWED CARTER'S LEAD BY GETTING 21850 VOTES TO HIS 20515.

Or if, hypothetically, Carter's current delegate count were now smaller than Kennedy's, the program should have inferred that

(j) IN THE PRIMARY ON 20 FEBRUARY, KENNEDY OVERTOOK CARTER BY GETTING 21850 VOTES TO HIS 20515.

instead. If we want good text from our generators, we have to give them the ability to recognize that "beat" or "lose" or "narrow lead" can be used instead of just the straightforward sentences of (f).

This ability is more than simply grouping together the two outcomes. It is an act of generator-directed inference, of interpretation, forming out of the two topics a new topic, perhaps one that does not even exist in memory yet. And the new topic is not simply a generator construct, but is a valid concept in memory. The act of determining that "beat" is appropriate is the act of interpreting the input as an instance of the concept BEAT denying this is to imply that "beat" can logically be used where BEAT is not appropriate, which is a contradiction.

A result of this view is the claim that lexical entities — words and phrases — can be accessed only via conceptual entities. (This may not be completely true, since analyses of certain kinds of slips of the tongue indicate that lexical items can also be accessed purely phonologically. However, phonological processes are not germane here, since they (presumably) occur only at a later stage in the generation of spoken language.) It is not clear how lexical entities would have to be organized to enable non-conceptual access. What other factors could possibly facilitate such access? How would one get "beat" from the two concepts Carter's and Kennedy's amounts of delegates, if not via semantics? This claim provides some comfort for lexicon builders. It implies that lexical entries need be associated only with their conceptual counterparts, in other words that additional inter-entry linkages are unnecessary. To claim that people could use the word "beat" without really understanding that they are expressing the concept BEAT is ridiculous.

This point has further ramifications. One could hold that the task of finding "beat" to satisfy a syntactic or pragmatic goal is a legitimate generator function, whereas the task of instantiating it as a concept and incorporating it into memory is not. However, it is clearly inefficient for the generator to interpret its input, say it, and then simply to forget

it again! — especially when there is no principled reason why the generator's inferences should be separated from other memory processes. Instead, after interpretation, the newly-built instance of the concept should be added to the system's representation of the story, where it can also be used by other processes or by the generator itself the next time it tells the story. In this way the content of memory can change as a result of generation. This explains our intuition that you often understand a topic better after you have told someone about it: the act of casting concepts and their interrelationships into coherent sentences has caused you to make explicit and remember some information you didn't have explicit before. In other words, not only does thinking influence talking, but also talking influences thinking.

Immediately, this view poses the question: *which process is responsible for making these inferences?* The possible positions on this issue reflect the amount of work one expects the generator to do. According to the strict minimalist position — a position held by most, if not all, generator builders today —, the generator's responsibility is to produce text that faithfully mirrors the input topics with minimal deviation: each input topic of sentence level produces a distinct output sentence (though perhaps conjoined with or subordinated to another). This minimalist position derives from a (presumably) unconscious reliance on linguistic arguments made by grammarians two decades ago. The arguments assumed (tacitly or otherwise; see for example [Chomsky 65, pp 148-163]) a separation between the processes that perform syntactic and semantic tasks. This separation is the antecedent of the belief of today's generator builders that tasks such as inference toward appropriate interpretations, being "semantic", are not properly the concern of a generator.

By separating out any tasks that operate upon or alter the semantic content of the domain, these grammarians and generator builders relegate the generator to its traditional position: that of passive back end, a more or less standalone module that can be given representation elements and then forgotten. This view cripples generators. Properly, a generator should be fully integrated in the main processing system, should have full access to the system's inferential capability, and should be able to activate processing goals and control the expansion of plans where appropriate. With this architecture, the above question doesn't arise: in the main system, the generator's inferential needs are indistinguishable

from those of any other subprocess.

3.3 Controlling Inferences

3.3.1 Unfortunate Practical Realities

Though linguists' arguments for the separation of syntactic and semantic processes might be motivated within a transformational approach to language, the arguments are certainly not appropriate for determining the position of generator builders. The result of holding such inflexible minimalist attitudes is that programs produce text that resembles the JUDGE examples (a) and (b). To circumvent this problem, in practice, most generator builders employ in their programs a number of special-purpose techniques, such as sophisticated sentence specialists that are sensitive to the subsequent input topics. Of course, this is a tacit acknowledgement that the minimalist position is not tenable.

In practice, unfortunately, generator programs are always back ends for other systems. The ideal conditions described above are never found. And, since generators can hardly expect these other systems to care about rhetorical and stylistic concerns, they have to perform interpretive inference under their own power. Thus, on renouncing the hard-line "no-inference" position, but not having access to the resources of the parent system, one is forced to face the question *how much inference must the generator do?*

I do not believe that a simple answer can be given to this question. The issue here is the same as that faced by planners in general: how much time can be spent developing alternative plans, checking their trustworthiness, and determining the likely costs and the likely benefits? In generator terms: a tradeoff exists between the time and effort required to make alternative formulations of the topic (which includes finding candidate interpretations, making them, and deciding on one) on the one hand, and the importance of flowing, good text on the other. Greater expense in time and effort produces better text. But, of course, the expenditure of these resources is controlled by the speaker's goals and other pragmatic

(i.e., interpersonal and situational) concerns. Thus such pragmatic criteria are appropriate for treating this question.

What are the likely benefits of running interpretation inferences? If they are successful, they may produce interpretations that:

- contain *fewer details* than the original topic(s) (as PAULINE's example (c) "Jim died in a fight" contains no details about individual blows). The relevant pragmatic questions are: does the hearer know (some of) the details already? Is he able to infer them himself? Even if he were able to infer them, should they be said explicitly in order to lend them more importance? These questions relate to Grice's maxim of manner (Grice 75), and are discussed below.
- are more *appropriately slanted*. Interpretations often add new facets or aspects to the topic; both the inclusion of new information and the exclusion of old details may slant the text. Here, relevant questions are: Can interpretations be found that state or imply the desired slant? Can interpretations be found that drop details working against the slant? These questions relate to Grice's maxim of quality; a discussion of strategies controlling opinion appears in chapter 4.
- contain some *additional information*. The inferences required to uncover new interpretations can, of course, provide additional information, even if it is only conjectural or default. For example, in a memory organized as PAULINE's is, interpretations can provide indices to similar instances as remainings. The information contained in these instances may play a pragmatic role as well, as discussed later in this chapter.

Balancing this, of course, is the fact that running inferences uses time. Thus the speaker must compare the importance of its goal to keep the conversation running against the importance of its goals to present an appropriate amount of detail to slant its text and to include additional information. PAULINE's answer to the question is *given the underlying pragmatic goal to express myself appropriately and well, I'd do as much inference as I can do, taking into account the available time, how much and what I want the hearer to know*.

and the richness of my memory and my lexicon. No doubt this answer offends people who like systems with well-defined boundaries; however, they must decide if they are happy with single-minded texts such as those above, and if not, they must provide an alternative without in any way doing interpretive inference for or otherwise. When the onus falls on the generator, built as the output end of some other system to produce good text, then it must be allowed to use as many means as are at its disposal.

3.3.2 Bottom-Up Inference

The problem in interpretive inference is to find valid interpretations easily and quickly. One solution to this problem is to try inferences directly on the input topics. This bottom-up method uses the structure of the memory network itself.

In PAULINE, bottom-up interpretive inference uses memory and the lexicon as part of the definitions of concept types. In order to find these inferences bottom-up, rather than have the pattern matcher check all patterns in memory, links are defined from concept types to the interpretations they can take part in. This scheme turns a concept representation network slightly different from the usual map-ports schemes such as, say, Stone & Brown 85, Clement, Roshok & McDermott 80, McDermott 80, and Brown & Winograd 77. Links are not defined to the interpreted concepts themselves, but to the information patterns. Note that the information patterns exist separately from the concepts that describe the causes and information, and the single interpretation is able to furnish more than one concept depending on other information. For example, BEAT and POSH can share information where the specific interpretation of a concept chosen may depend on the symptoms of the generation thus in which it wants to interpret.

Still more fully interlinked network the program now has to manage a number of trackless inferences. In order to limit further the number of interpretations PAULINE has to manage, only some concepts are linked to the information as they are part of. Typically, these are the more central ones, the ones linked to most others. These central concepts are

called *pivot concepts*. For example, PAULINE can access the configuration for NARROW LEAD from an outcome or from the main primary election, but not directly from the current delegate counts or from the relations between concepts. This strategy can cut down on the search time considerably, but it also means that PAULINE is unable to find an interpretation for which it has all but the pivots as input topics. For this reason, when it has on, with time (that is, when its rhetorical goals that control the level of detail and the amount of time permit), the program uses the aspects of its current input topics, as well since they might be pivots even when the topics themselves are not. Still, for a large set of input topics containing many pivots, collecting all configurations can be a daunting task. Some further strategies to pare it down are, for example, selecting only one configuration from each pivot (or perhaps two, depending on the amount of time available), or using only a subset of the input topics (determined, by, for example, the generator's sympathies).

Of course, this is not a wonderful inference system — it depends on the right links being defined beforehand — but it is an acceptable solution for limited domains. The implementation of a full-blown inferencing scheme is not a generation problem. Whether you define links from concepts to possible configurations or associate with concepts appropriate inferences, you are simply simulating the action of a mechanism that provides you with a number of (hopefully appropriate) candidate interpretations in a reasonably short time.

3.3.3 Top-Down Inference

Another solution is to run only inferences that are likely to produce useful results. But where can such inferences be found? One source is the plans that serve the generator's goals. Potentially useful interpretation inferences can be explicitly included in these plans, in such a way that running a plan causes appropriate inferences to be applied to the collected candidate sentence topics. Since interpretation is such a powerful way of slanting the text, the generator's rhetorical goals of opinion are an eminently suitable source of guidance. Indeed, many of these goals can *only* be achieved through performing appropriate interpretation inferences on the input topics.

In the shantytown examples, a number of top-down inferences were used. These inferences were obtained from the rhetorical goals of opinion, described in chapter 2:

- **Present as confrontation:** state that the actor you oppose (B) did some action (ACT) as a confrontation with some actor you support (A). In more detail, this rule can be represented as:

```
IF B has the goal that some actor X must do some action Z
AND A has goal that X must do Z'
AND Z' conflicts with Z
AND B's action ACT forces X to do Z' (disregarding A)
THEN interpret and present ACT as a confrontation with A
```

- **Present as conciliation:** state that the actor you support (A) did some action (ACT) in the spirit of conciliation. Represented as:

```
IF A has some goal G
AND B has some opposing goal G'
AND ACT serves some other goal H
AND H does not conflict with G'
AND H is (or serves) a goal of the opponent B
THEN interpret and present ACT as an act of conciliation
```

- **Present as coercion:** state that the actor you oppose (B) forces someone else (X) to do some action (ACT). Represented as:

```
IF ACT serves one of B's goals G
AND G opposes the goal(s) of the actor you support
AND X, the actor, is not directly an agent of B
AND B has had interactions with X
```

THEN interpret and present B's action as coercing X into
doing ACT (and use 'force' and additional verbs
such as 'have do', 'cause')

- Present as appropriation: state that the actor's purpose B is to take something
into the ACT. Represent as

IF ACT serves one of B's goals
AND and instrumental to ACT,
the actor used some thing INSTR
AND INSTR does not belong to the actor or to B,
but to someone else
THEN interpret and present the use of INSTR as an
appropriation (and use 'take over', 'grab')

- Present as deceit: state that the actor's purpose B is to say something that is not
true. Represent as

IF ALL serves one of B's goals
AND it is the transfer of a message
AND some aspect of the message conflicts
with veridical information in the world
THEN interpret and present ALL as a well-falsified

Interpretation of the above rules is that they are intended to be used in a way that is consistent with the following principles of interpretation: (1) The actor's purpose is to be interpreted as the purpose that is most likely to be the cause of the action. (2) The actor's purpose is to be interpreted as the purpose that is most likely to be the cause of the action.

Interpretation of the above rules is that they are intended to be used in a way that is consistent with the following principles of interpretation: (1) The actor's purpose is to be interpreted as the purpose that is most likely to be the cause of the action. (2) The actor's purpose is to be interpreted as the purpose that is most likely to be the cause of the action.

IN EARLY APRIL, A SMALL NUMBER OF STUDENTS [WERE INVOLVED IN A CONFRONTATION] WITH YALE UNIVERSITY OVER YALE'S INVESTMENT IN COMPANIES DOING BUSINESS IN SOUTH AFRICA. THE STUDENTS [TOOK OVER] BEINECKE PLAZA AND CONSTRUCTED A SHANTYTOWN NAMED WINNIE MANDELA CITY [IN ORDER TO FORCE] THE UNIVERSITY TO DIVEST FROM THOSE COMPANIES. YALE REQUESTED THAT THE STUDENTS ERECT IT ELSEWHERE, BUT THEY REFUSED TO LEAVE. LATER, AT 5:30 AM ON APRIL 14, OFFICIALS HAD TO DISASSEMBLE THE SHANTYTOWN. FINALLY, YALE, [BEING CONCILIATORY] TOWARD THE STUDENTS, NOT ONLY PERMITTED THEM TO RECONSTRUCT IT, BUT ALSO ANNOUNCED THAT A COMMISSION WOULD GO TO SOUTH AFRICA IN JULY TO EXAMINE THE SYSTEM OF APARTHEID.

I AM SAD TO SAY THAT A FEW STUDENTS [TOOK OVER] BEINECKE PLAZA IN EARLY APRIL AND BUILT A SHANTYTOWN CALLED WINNIE MANDELA CITY. THEY WANT YALE UNIVERSITY TO DIVEST FROM COMPANIES DOING BUSINESS IN SOUTH AFRICA. OFFICIALS HAD TO TEAR IT DOWN, BECAUSE YALE WANTED THINGS TO BE ORDERLY. FINALLY, THE UNIVERSITY [REVERSED] AND ALLOWED THE STUDENTS TO PUT IT UP AGAIN. YALE'S DESIRE TO BE REASONABLE WAS A GOOD THING.

PAULINE: [The students' confrontation with the university appears to be a confrontation with the university's investment in South Africa, which is a confrontation with the university's investment in South Africa.]

3.4 How PAULINE Does It

PAULINE: [The students' confrontation with the university appears to be a confrontation with the university's investment in South Africa, which is a confrontation with the university's investment in South Africa.]

or patterns of concept types and relations among them. The interpretation mechanism matches candidate configurations against its collected topics, and, if matched, creates a new instance of the interpretation and adds it to memory. The program can then generate text from the interpretation instead. Details of the process appear in chapter 7.

Configurations: In order to find applicable interpretations, PAULINE uses patterns called configurations. A configuration is the description of the way in which a collection of concepts must relate to each other to form a legitimate instance of a high-level concept. As described in chapter 7, a pattern-matcher matches the candidate topics against likely interpretation configurations.

Matching configurations: During the topic organization stage, PAULINE gathers likely interpretation inferences (both top-down and bottom-up) and, using a simple pattern-matcher, applies their configurations to the set of collected candidate topics and collects all the matched occurrences. When it has a partial match — partial in the sense that the remaining concepts are not among the candidate topics — the program can either accept or reject the match. In the latter case, it can use the relationships in the configuration to search memory for the remaining concepts, though they may not originally have been meant as topics. At present, PAULINE does not do this (it simply rejects partial matches); adding this capability to the program would be a simple extension, requiring in addition a pragmatic-based decision that could be characterized as, say, **thoroughness**, and perhaps be implemented as a rhetorical goal with the possible values *pedantic*, *neutral*, *lax*.

Selecting configurations: When a number of configurations have been matched, either through top-down control or through bottom-up methods, the generator must select which ones to say. Three pragmatic factors play a role: **interest**, **affect** (slant), and **reminding**.

(a) **Interest:** This relates to the number of concepts contained in a configuration. With respect to interest, interpretations provide a way to compress many concepts and say them briefly so as not to bore the hearer. However, the process of finding a suitable interpretation can take some time and is not even guaranteed to work, so as a time-saving

strategy the use of interpretations is not reliable.) Appropriate strategies depend on the rhetorical goal **RG:detail** (which is discussed below):

- select the largest configuration (i.e., be most concise) when **RG:detail** is *interpretations*
- say no configuration (i.e., say all the details) when **RG:detail** is *detailed*
- select an intermediate configuration otherwise

(b) **Affect:** This relates to the number of affectively sensitive concepts in a configuration. With respect to affect, interpretations provide a way of including topics into the conversation without actually saying them explicitly, thereby satisfying both the speaker's need to say them and the hearer's need not to hear them. For instance, in the Carter-Kennedy example, by choosing to say

(1) IN THE PRIMARY ON 20 FEBRUARY, KENNEDY NARROWED CARTER'S LEAD

the speaker avoids explicitly mentioning Carter's defeat. These strategies are independent from those of interest, since a large configuration may contain fewer or more sensitive concepts than its smaller alternatives (sets of small configurations seldom span the same concepts as large ones). The strategies depend on **RG:partiality** (which is discussed in chapter 4).

- select the configuration with most sensitive concepts when **RG:partiality** is *partial*
- select the configuration with fewest sensitive concepts when **RG:partiality** is *impartial*
- otherwise, select a configuration with some intermediate number of sensitive concepts

(c) **Reminders:** This relates to the presence of other instances in memory similar to the interpretations, and is discussed later in this chapter.

Making interpretations: When a group of concepts matches the pattern of a configuration, this group can be interpreted as an instance of the configuration's associated concept. If an interpretation with the group's concepts is not already present in memory, a new high-level interpretation can be created and indexed off the interpretation, following the memory organization principles described in [Schank 82]. PAULINE creates a new instance of the interpretation type and fills it with the concepts from the group that it found had matched the configuration. It then places it in memory and links the pivot concept(s) to the configuration that matched.

For example, the first time PAULINE generates a Carter-Kennedy story, it creates the new interpretations BEAT and NARROW-LEAD and adds them to the story representation. The next time PAULINE generates the story, it finds the two new interpretations immediately, using links from *one of the pivot concepts*, thereby avoiding the search and matching process. (Of course, at this point, the program tries to make further interpretations off these, but finds no appropriate concepts in its limited memory.) Thus, as a result of having said this once, memory has been extended, and PAULINE can be said to "understand" the topic better.

Generating interpretations: Finally, the interpretation can be said. Since it is a standard concept (not, say, a generator-specific construct), generation proceeds normally. Of course, the interpretation replaces the topics it subsumes; for example, after deciding to say BEAT or NARROW-LEAD, the input topics Carter's and Kennedy's outcomes become redundant.

3.5 Determining the Appropriate Level of Detail

As described in chapter 2, different textual aspects are controlled by different speaker goals. In this chapter, we have been discussing the level of topic detail. A final question remains: what is the appropriate level of detail in any given circumstance? That is, when does the hearer need to know the details of the topic? What is the effect of telling him only high-level interpretations? Or of telling him both? The answer can be summarized as: if you can

trust him to make the high-level interpretations himself, then all you need give him are the details. To repeat the two relevant sentences from the Carter-Kennedy example:

(f) IN THE PRIMARY ON 20 FEBRUARY, CARTER GOT 20515 VOTES.
KENNEDY GOT 21850.

(i) IN THE PRIMARY ON 20 FEBRUARY, KENNEDY NARROWED CARTER'S
LEAD BY 1335 VOTES.

If the hearer is a political pundit and he is following the nomination race with interest, then clearly (f) is better, since he can draw the conclusion without difficulty, and, in addition, he has precise numerical information. If, in contrast, the hearer has only minimal knowledge about or interest in the nomination procedure, then (i) is better, since it doesn't burden him with details and require him to do the interpretation himself. What must you say, however, if the hearer is interested and has a limited amount of knowledge — say, he is a student of the political process —, or if he is knowledgeable but unlikely to make the right interpretation — say, he is a strong Kennedy supporter, whereas you are pro-Carter? In both these cases you must ensure that the hearer understands how you expect him to interpret the facts. So you tell him details *and* the interpretations:

(m) KENNEDY NARROWED CARTER'S LEAD IN THE PRIMARY ON 20 FEBRUARY.
HE GOT 21850 VOTES AND CARTER GOT 20515.

These considerations can be stated as the following rules (using the terms defined in chapter 2 to characterize the pragmatic aspects of conversations and goals of speakers). PAULINE uses these rules to activate the rhetorical goal **RG:detail** that controls the level of detail of topics generated. The goal takes one of the values *details*, *interpretations*, *all* or *no details and interpretations*):

- set **RG:detail** to *details* if the hearer is likely understand the details or if he wants to hear the details. This rule bears on his background knowledge, and in **PAULINE**, it is decided by referring to its information about the hearer: is the **hearer's knowledge level** marked *expert* (does he know enough about the topic to be able to understand the details and their significance?); or is the **hearer's interest level** marked *high* (does he not care enough about the answer to want to hear anything but an interpretation?)
- otherwise, set **RG:detail** to *all* if he is likely to make the wrong interpretations of the details. This rule depends on various factors: is the **hearer's knowledge level** marked *student* or *novice* (does he have too little inferential knowledge to be able to make the interpretation?); or is the **atmosphere (time)** not marked *little*; and finally, will different sympathies cause him to make a different interpretation? (check the hearer's sympathies and antipathies for the central topic of the conversation)
- otherwise, set **RG:detail** to *interpretations*

In addition to these considerations, the value of the goal can be affected by the desire not to upset the hearer's sympathies:

- then, set **RG:detail** to *interpretations* if it is better to avoid painful topics, to ensure that painful aspects (the details, the interpretation, or the inferences used to make it) can simply be left out. This rule translates as follows: is **speaker-hearer depth of acquaintance** marked *strangers*, or is **speaker-hearer relative social status** marked *subordinate*, or is **desired effect on hearer's emotion toward speaker** marked *like*, or is **desired effect on interpersonal distance** marked *close*, or is **desired effect on hearer's emotional state** marked *calm*?

In summary, you must be as specific as the hearer's knowledge of the topic allows: if you are too specific he won't understand, and if you are too general you run the risk of seeming to hide things from him, or of being uncooperative. In the first case, you violate the default speaker goal to be intelligible, and in the second, you violate the goal to avoid

unacceptable implications. In either case, you violate Grice's maxim of quantity to say neither more nor less than is required ([Grice 75]).

3.6 The Inclusion of Reminders

3.6.1 Reminders

In [Schank 82], Schank describes how memory is organized so that specific instances are indexed off general concepts to aid generalization and explanation. These specific instances come up during processing as reminders. Since a generator's interpretations are themselves concepts, they can furnish reminders; these reminders can be used as examples in argumentation. That is to say, if the generator has the goal to say a number of concepts, and it finds a interpretation which neatly expresses the concepts and their relations, and the interpretation can furnish a specific instance of itself, this instance will be relevant and can be used to strengthen the argument.

In the Carter-Kennedy example, the concept NARROW-LEAD was provided with two instances: the instance when Carter narrowed Udall's lead in a primary in 1976, and the time when Hart narrowed Mondale's lead in 1984. (In a fictitious world such as PAULINE's, anachronistic reminders are no stranger than normal ones! The names Hart, Mondale, and Udall were simply chosen because similar instances did in fact occur to them during their bids for the nomination.) When biased, PAULINE uses an appropriate reminding:

(n) KENNEDY DIMINISHED CARTER'S LEAD BY GETTING ALL OF 21850 VOTES IN THE PRIMARY IN MICHIGAN. IN A SIMILAR CASE, CARTER DECREASED UDALL'S LEAD IN A PRIMARY IN 1976, AND HE TROUNCED UDALL TO BE NOMINATED BY 2600 DELEGATES. I AM GLAD THAT KENNEDY IS NOW CLOSER TO GETTING THE NOMINATION THAN HE WAS BEFORE.

(o) KENNEDY SLIGHTLY DIMINISHED CARTER'S LEAD IN THE PRIMARY IN MICHIGAN. IN A SIMILAR CASE, HART DECREASED MONDALE'S LEAD IN

1984, BUT MONDALE EASILY BEAT HART TO BE NOMINATED BY 1500 DELEGATES. CARTER STILL HAS MANY MORE DELEGATES THAN KENNEDY DOES.

A reminding found off an interpretation concept obviously has the same structure as the configuration of the relevant input topics. That is what makes it relevant. The details, of course, are different. That is what makes it interesting. Sometimes a reminding may contain parts that do not correspond to any input topic, requiring the generator to decide whether these parts should be included. (There is, for example, no equivalent to Carter's beating Udall for the nomination, since the Carter-Kennedy nomination has not yet, in the hypothetical example, taken place.) Sometimes a concept may furnish more than one reminding, and the generator has to pick the one relevant to its goals. Just as it did when selecting an interpretation, the generator must choose a reminding whose details serve its goals.

In order to select a reminding, a mapping must be set up between the input topics and/or interpretation and the reminding, so that the corresponding actors, objects, times, places, etc., are determined. Aspects of the reminding must be replaced with aspects of the input to create the hypothetical case, which can then be used to determine affective suitability. Thus, in the examples, PAULINE "translates" the Carter-Kennedy case into both the Carter-Udall and the Hart-Mondale scenarios and finds that in the former, Kennedy would win the nomination (since in 1972 front-runner Udall lost), and in the latter, Carter would win (since in 1984 front-runner Mondale won). Depending on its sympathies, it selects a suitable reminding. At this point, the generator can either spawn the goal to say the reminding immediately, or it can start doing further topic collection from the aspects of the reminding. (In PAULINE, the decision is based on the rhetorical goal **RG:haste**: the less hasty, the more time to do further topic collection.) If said immediately, the reminding is woven into the text just after the concept that gave rise to it, using phrases such as "that reminds me" or "in a similar case".

3.6.2 Adding Color to the Text

Color means any reference to personal experience to illustrate general statements: specific instances (as reminders), idioms and frozen phrases (as interpretations of the situation), and descriptions of occurrences of personal evaluation.

When and why should the speaker include examples and idioms? What conversational circumstances and speaker goals influence the relevant rhetorical goal? When used appropriately, a well-chosen example makes abstract points clear and dull text alive. Nobody can forget the agony of suffering through a textbook containing few or no examples. Though the speaker seldom *has to* include instances and idioms, they certainly make the text more understandable. Also, as shown in (l) and (m), the speaker can strengthen the force of his argument by citing an appropriate example. In addition, since instances and statements of personal evaluation are in some sense interjections, they are the perfect way of repairing goals that are in danger of being thwarted by the text or even simply goals that have been ignored for too long. For example, why would the speaker refer to his French in

“You don’t speak proper English -- it’s better than my French, but it’s still not good — ...”

thereby doubling the length of the sentence? Clearly, because the speaker had active both the goal to express his evaluation of the hearer’s English and the goal to have the hearer like him. The former goal gave rise to the first part of the sentence. After it was said, the goal tracking mechanism inferred that that part of the sentence could be taken as an insult, which conflicted with the latter goal; so it had to be repaired. This conflict gave rise to the goal to mitigate the effects of the sentence by using *one of the following strategies*:

- say something good about the hearer: “ — but at least you write tolerably well”
- say something bad about the speaker: “ — but it’s better than my French”
- say how nobody else is any better: “ — though second languages are never perfect”
- motivate or help the hearer to improve: “ — how about attending a writing class?”

Usually, then, examples help to clarify the topic or help to serve other goals that require some attention. These are the rules PAULINE uses to establish a value (one of *with examples*, *normal*, *no examples*) for its rhetorical goal **RG:color**:

- set **RG:color** to *with examples* if the following goals are present: **desired effect on hearer's knowledge** is marked *teach* (since a paragraph explaining something is usually more effective if it contains an example; or if **desired effect on hearer's goals** is marked *activate* (for example, if the speaker has the goal to suggest to the hearer possible future plans and actions, since concrete examples are more direct and effective than simple injunctions; compare

"If you want to become rich, try the lottery"

"If you want to become rich, try the lottery. Last week a woman won 10 million when she used her family's birthdates").

Also, if **desired effect on hearer's emotion toward speaker** is marked *respect* or *like* and the speaker's and hearer's affects for the topic differ (since explaining his reasoning or reactions can help to make the speaker be better understood); or if **topic collection goal** is marked *convince* (that is, when the speaker wants to present support for his interpretation or affect; for example, in the following sentences,

- "Stalin, a ruthless man, was a charmer"

- "Stalin, a man who killed 15 thousand people, was a charmer"

the latter underscores the speaker's antipathy. Almost any concrete facts that are mustered in support of an argument lend it force).

In addition, **repair a goal that has been slighted** (by saying appropriate examples using one of the four strategies mentioned above)

- set **RG:color** to *no examples* if the following goals are present (since an example can aid most goals, it is perhaps best to note when they should *not* be included): **desired effect on interpersonal distance** is marked *distant*; or if **atmosphere (tone)** is marked *formal* (especially, for example, if the topic itself is embarrassing or intimate one does not describe your dental problems in a speech, even if appropriate); or if **desired effect on hearer's emotion toward speaker** is marked *dislike*; or if

desired effect on hearer's knowledge is marked *confuse* (that is, when the speaker doesn't want to be understood)

The following strategies are associated with **RG:color** and serve to increase the color of text by selecting appropriate options at the following decisions points:

- **topic selection and inclusion:** include, as examples, other instances similar to the topic, such as those found off interpretations
- **topic inclusion:** summarize an argument or a point by including an appropriate idiom rather than general statements, for example by adding "So don't count your chickens before they're hatched!" to texts (n) or (o) (PAULINE cannot add such phrases, though it clearly has some of the requisite information and goals)
- **topic inclusion:** include sentences describing personal evaluations
- **clause inclusion:** make adjectival clauses of appropriate instances, for example the Stalin example above
- **phrase/word selection:** select metaphoric and idiomatic phrases and words, such as "crowned with the nomination" rather than "got the nomination"

3.7 Conclusion

As generators become larger and more complex, and as they are increasingly used together with other programs, they should use the capabilities of those programs to further their own ends, and, especially, to produce better text. Therefore, we should study the kinds of tasks that generators share with other processes and the purposes generators require them to fulfill. This chapter describes some of the kinds of demands a generator can be expected to place on a general-purpose inference engine. And even with the limited inferential capability described here, PAULINE can greatly enhance the quality of its text and the efficiency of its communication of non-literal pragmatic information.

Chapter 4

Affect in Text

Abstract

This chapter discusses the communication of opinion. Natural languages contain a large number of linguistic techniques for slanting text — techniques that control both what to say and how to say it. All decisions made by these techniques are based upon one general rule, the affect rule, which is derived from the goals all speakers must have in order to ensure their hearers' attention. In order to use the techniques, a generator must be given opinions and the ability to derive opinions for related topics.

4.1 Introduction

Any speaker who is sensitive to the pragmatic aspects of conversation must be able to include his opinions in his text. People do this all the time; our biases sneak into what we say so easily and so often that producing genuinely unslanted text can be quite a problem! And when we do manage it, the resulting text is often boring. For example, the following text in neutral newspaper style:

(a) YALE UNIVERSITY PUNISHED A NUMBER OF STUDENTS FOR BUILDING A SHANTYTOWN, WINNIE MANDELA CITY, ON BEINECKE PLAZA BY ARRESTING 76 STUDENTS AND TEARING IT DOWN ONE MORNING IN APRIL. THE STUDENTS WANTED YALE TO DIVEST FROM COMPANIES DOING BUSINESS IN SOUTH AFRICA. FINALLY, THE UNIVERSITY COMPROMISED AND ALLOWED THE STUDENTS TO REBUILD IT.

certainly makes less stirring reading than the slanted texts:

(b) I AM ANGRY ABOUT YALE'S ACTIONS. THE UNIVERSITY HAD OFFICIALS DESTROY A SHANTYTOWN CALLED WINNIE MANDELA CITY ON BEINECKE PLAZA AT 5:30 AM ON APRIL 14. A LOT OF CONCERNED STUDENTS BUILT IT IN EARLY APRIL. NOT ONLY DID YALE HAVE OFFICIALS DESTROY IT, BUT THE POLICE ARRESTED 76 STUDENTS. AFTER THE LOCAL COMMUNITY'S HUGE OUTCRY, THE UNIVERSITY ALLOWED THE STUDENTS TO PUT THE SHANTYTOWN UP THERE AGAIN.

(c) IT PISSES ME OFF THAT A FEW SHIFTLESS STUDENTS WERE OUT TO MAKE TROUBLE ON BEINECKE PLAZA ONE DAY --- THEY BUILT A SHANTYTOWN, WINNIE MANDELA CITY, BECAUSE THEY WANTED YALE UNIVERSITY TO PULL THEIR MONEY OUT OF COMPANIES WITH BUSINESS IN SOUTH AFRICA. OFFICIALS HAD TO CLEAR IT FROM THERE. FINALLY, YALE GAVE IN AND LET THE SHITHEADS PUT IT UP AGAIN, AND YALE SAID THAT A COMMISSION WOULD GO TO SOUTH AFRICA TO CHECK OUT THE SYSTEM OF APARTHEID.

Clearly, in any real account of the episode, the protesters' version is going to differ appreciably from the university's. The differences will not be haphazard; each speaker will make the decisions that slant the text in his favor. What are these decisions?

The previous chapter describes how one of the possible effects of interpreting input topics is the slanting of text. But, as these examples show, that is by no means the only

technique to use. Note the adjectives "concerned" and "shiftless (students)" and "huge (outry)": the statements of explicit opinion, such as "I am angry" and "it pisses me off"; and the suppression of topics that are inappropriate to each point of view, such as the investigation commission in (b) and the local community's outcry in (c). Biases can be stated explicitly, using statements such as "I think X is good", or can be injected into the text at various points during the realization process. The latter is often more effective, since it incorporates opinion implicitly into both the content and the form of sentences.

PAULINE uses these strategies to produce the following two versions of the representation of a fight (built by the JUDGE program, a case-based expert system that models the sentencing behavior of a judge, as described in chapter 1; see [Bain 86, 84]). If PAULINE's sympathies are for Mike, the program says:

(d) FIRST, JIM BUMPED MIKE HARD AND HURT HIM.

MIKE JUST TAPPED JIM ONCE.

AFTER THAT, JIM DID NOT EXPECT THAT MIKE WAS GOING TO HURT JIM 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. THE ACTION WAS A SIMPLE RETALIATION.

NEXT, MIKE HIT JIM, KNOCKING HIM DOWN.

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.

MIKE REALIZED THAT JIM REALLY HURT MIKE; ALSO, MIKE EXPECTED THAT JIM'S GOAL WAS TO HURT HIM. JIM WAS STABBED.

AS A RESULT, JIM ONLY DIED.

(...taking mitigation to its (il)logical extreme!). From the same input, if PAULINE is defending Jim, it stresses Mike's actions and culpability and minimizes Jim's:

(c) 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 JIM WOULD HURT HIM. HE HAD NO INTENTION TO BOTHER MIKE. MIKE ONLY WAS BUMPED BY JIM ONCE. THE ACTION WAS AN ACCIDENT.

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.

NEXT, JIM REALIZED THAT MIKE INJURED JIM. JIM JUST TAPPED MIKE ONCE. THE ACTION WAS A SIMPLE RETALIATION.

MIKE DID NOT FORESEE THAT JIM WAS GOING TO HURT HIM ANY LONGER. MIKE WANTED JIM TO FALL. MIKE COULD EXPECT THAT JIM WOULD BE KNOCKED TO THE GROUND BY MIKE IF MIKE HIT JIM HARD ONCE. HE SMASHED HIM, KNOCKING HIM DOWN. HE WAS NOT JUSTIFIED FOR KNOCKING HIM TO THE GROUND.

JIM REALIZED THAT HE FELL. NEXT, HE ONLY SLAPPED MIKE.

AFTER THAT, MIKE FORESAW THAT JIM WANTED TO HURT HIM. MIKE WANTED TO KILL JIM TO PREVENT HIM FROM SLAPPING MIKE. HE COULD EXPECT THAT IF HE STABBED JIM SEVERAL TIMES HE WOULD KILL HIM. HE STABBED HIM. THE RESULT WAS THAT HE WOUNDED HIM. MIKE'S CAUSING JIM TO BE KILLED WAS NOT JUSTIFIED.

AS A RESULT, JIM DIED.

In this chapter, two additional sets of techniques for slanting text are described: *content-related* slanting techniques, involving judicious topic selection, and *form-related* techniques, involving the use of enhancers and mitigators, the appropriate juxtaposition of topics in

phrases, the ordering of sentence parts, and choices of subject, clause content, and words. First, however, we must discuss how to give a program opinions.

4.2 Computing Affect

What, now, is the meaning of "slanting the text in Jim's favor"? Clearly, to present someone favorably is to try to get the hearer to view him sympathetically. In order to do this, the speaker must be able to distinguish between what the hearer is likely to find sympathetic, what he is likely to dislike, and what he is likely not to care about much. Therefore, three values of affect are required to model this decision: GOOD, BAD, and NEUTRAL. Of course, human affect is not a simple value on a linear range. But *affect* is not used here to mean all human feelings. It simply denotes something akin to "like". And with this limited denotation, three values are sufficient to give the program a great deal of interesting behavior. In this regard it is similar to the work on narrative summarization done by Lehnert, which also uses a simple three-valued affect with very interesting results ([Lehnert 82]).

In general, affect derives from two sources: from the speaker's opinions about some contentious topic, and from the default intrinsic affects associated with concepts. To a generator program, these sources are the user and the intrinsic affects defined for the representation elements. For PAULINE, the first source is simply implemented by having a sympathy and an antipathy list. Elements on these lists will be characterized as GOOD and BAD respectively. In the JUDGE examples, when PAULINE defends Mike, the sympathy list contains the concept "Mike" and the antipathy list the concept "Jim". In the Carter-Kennedy examples, PAULINE is made a Carter supporter by marking the element representing Carter's goal to win the nomination as GOOD and marking Kennedy's goal as BAD. In the shantytown examples, PAULINE is given one protagonist and his goals as sympathies and the other and his goals as antipathies.

The second source of affect is tied to the generic representation elements. Each representation type that carries some intrinsic affect in the example domain has this affect

defined. 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, the given affects must be combined with the concepts' intrinsic affects and must be propagated along the relations between concepts. The basic rules are:

Rules of affect propagation:

- affect is preserved when combined with NEUTRAL
- like affects combine to GOOD
- unlike affects combine to BAD
- when the two affect-bearing concepts are related in certain ways the combined affect inverts (for example, when one affect bearer is a BAD action and the other is its conceptual patient). Rules of affect propagation must be defined for each possible relationship between affect-bearing concepts

although their exact form obviously depends on the design of the representation. This rule is similar to the "balance principle" discussed in [Abelson & Rosenberg 58]. A description of how PAULINE computes affect from input representations is given in chapter 7.

Which affect the hearer is likely to have for a representation element is computed in exactly the same way, except that the initial sympathies and antipathies are taken from the hearer's sympathies and antipathies.

4.3 The Affect Rule and its Application

Knowing what affect the hearer has for each piece of the representation does not yet tell the speaker what to do. He requires strategies that indicate when and how to say GOOD

things and when and how to say BAD ones. With regard to affect, the strategies used in most conversations are all based upon one very general rule, which is called here the *affect rule*.

All speakers have the low-level goal to ensure that the hearer will be receptive to the implications of the generated text. That is, the speaker must not say things that will offend the hearer and cause him to terminate the conversation. Differences of opinion — i.e., conflicting affective values for topics — is a common cause of offense. Since his sympathies and antipathies reflect so accurately the speaker's disposition toward the world, any opinion with which the hearer disagrees signals distance between them, and perhaps even censure on the part of the speaker. The speaker should avoid such opinions if he wants the hearer to accept his conversation. Therefore, this goal requires that, whenever there is a conflict, or even a potential conflict, between the affective implications of the text and the hearer's opinion, the speaker has to skirt sensitive issues and achieve effects indirectly. On the other hand, when the speaker and the hearer agree on the affect of a topic, this agreement can be stressed to emphasize their agreement.

Fortunately, most languages have a large body of techniques for skirting or emphasizing issues. The application of these techniques (where *enhancers* are words or phrases that strengthen the affect of a concept, and *mitigators* are words or phrases with the opposite effect) is controlled by the

Affect Rule:

To convince the hearer that some topic is GOOD or BAD, combine it with other GOOD or BAD topics using enhancers and mitigators:


- for a GOOD effect, say GOOD topics with ENHANCERS and BAD topics with MITIGATORS
- for a BAD effect, say GOOD topics with MITIGATORS and BAD topics with ENHANCERS


AD-A183 418 GENERATING NATURAL LANGUAGE UNDER PRAGMATIC CONSTRAINTS 2/3
(U) YALE UNIV NEW HAVEN CT DEPT OF COMPUTER SCIENCE
E H HOVY MAR 87 YALEU/CSD/RR-521 N00014-85-K-0108

UNCLASSIFIED

F/G 5/7

NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

This rule is the basis for all affect-related generation decisions. (It is, however, clearly not applicable to all conversations. In an argument, the speaker may violate the affect rule and still have the hearer listen to him. In this case, the hearer is receiving alternative forms of “payment” — even if the only reason he continues the argument is his reluctance to walk away and appear the loser. Another common violation is a comedy act in which the comedian insults his audience. Here the alternative payment is entertainment; if the comedian isn’t funny, he loses. All such cases can be described as exceptions to the general rule, and their strategies as inversions of the normal strategies based on this rule. The affect rule is the basis on which everything else rests.)

The techniques PAULINE uses to enhance and to mitigate topics are described next. The use of these techniques is determined by strategies that are applied to the options at decision points during the generation of sentences, in exactly the same way that the strategies to achieve the rhetorical goals are applied.

4.4 Content: Topic Collection Techniques

Part of a generator’s task is to determine what to say. Though it is possible to claim that the generator should say only what it is given, in other words that the task of finding and filtering topics belongs to some other process, this constraint can hamper the production of pragmatically sensitive text, of affect-laden text in particular. For assume the “other process” decides it is important to say that that Jim was stabbed by Mike. Then if the generator has the goal to support Mike, it should legitimately be able to decide not to say that at all, or at least to mitigate the bald statement “Mike stabbed Jim” by, for example, “...but Jim bumped Mike first”. If it is not able to suppress topics or to find mitigating circumstances such as Jim’s starting the fight, it really cannot do much to help Mike.

By the affect rule, there is no problem when the speaker and hearer agree on the affect of the topic. Such topics can simply be said directly. But when they disagree, the speaker has to be more careful.

4.4.1 Evasion

One strategy the speaker can try is *evasion*: dealing with the topic indirectly, through hints and implications, by referring to something that is in some way related, and trust that the hearer will perform the final bridging inference himself. Various evasive techniques exist. The simplest one is:

- **Wishful Suppression and Mitigation:**

- say GOOD topics
- juxtapose NEUTRAL topics with GOOD ones in enhancer phrases
- leave out BAD topics altogether, unless they can be mitigated using mitigator phrases and words, or unless they are central to the story

PAULINE uses this strategy in the JUDGE examples. The input from the JUDGE program consists of a list of interpretations, where each interpretation describes an action, its justifiability, and the actor's motivations and culpability, in the JUDGE's opinion. To illustrate, a literal rendering (generated by the JUDGE program) of its interpretation for the first action of the fight is

(f) JIM KNEW THAT MIKE HAD NOT YET HURT HIM. HE COULD PERCEIVE NO THREAT AGAINST HIM FROM MIKE. A REASONABLE PERSON IN SIMILAR CIRCUMSTANCES WOULD HAVE PERCEIVED NO THREAT AGAINST HIM FROM MIKE. JIM COULD NOT FORESEE THAT MIKE WOULD BE HURT IF JIM ACCIDENTALLY BUMPED MIKE. A REASONABLE PERSON COULD NOT FORESEE THAT IF HE BUMPED HIM JIM WOULD HURT HIM. HE HAD NO INTENTION TO BOTHER MIKE. JIM BUMPED MIKE. THE RESULT OF THIS ACTION WAS THAT MIKE FELL. JIM'S INTENTION TO BUMP MIKE WAS NOT JUSTIFIED. IT IS SIMPLY UNCLEAR FROM THE INPUT AS TO WHY JIM ACTED AS HE DID.

Instead, when PAULINE is defending Jim, it computes its affect for each part of the interpretation and, using the **wishful suppression and mitigation** strategy, decides whether or not to include it. From Jim's point of view, almost all the parts of the interpretation are GOOD; after all, Jim didn't intend to hurt Mike. The first three sentence topics are NEUTRAL and are ignored. The fourth is GOOD; it is paired with the NEUTRAL fifth using the *also* enhancer. The sixth is GOOD, and the seventh must be included, since, in this domain, actions are defined to be the central elements of the story. The last sentence is also GOOD. Thus in defense of Jim the result is

(g) JIM, A NICE FELLOW, COULD NOT EXPECT THAT THAT UNPLEASANT MIKE WOULD BE HURT IF JIM ACCIDENTALLY BUMPED HIM; ALSO, A REASONABLE PERSON COULD NOT FORESEE THAT IF HE BUMPED HIM JIM WOULD HURT HIM. HE HAD NO INTENTION TO BOTHER MIKE. MIKE ONLY WAS BUMPED BY JIM ONCE. THE ACTION WAS AN ACCIDENT.

On the other hand, when PAULINE is defending Mike, there is much less to say. From Mike's point of view, the only GOOD part of the input is Jim's action and its result, and PAULINE must make the most of it. Therefore, out of all this, PAULINE chooses to say only:

(h) THAT NASTY FELLOW JIM BUMPED MY FRIEND MIKE HARD AND HURT HIM.

For the full effect of this strategy, compare the contents of texts (d) and (e) above.

Two other evasive techniques (neither of which PAULINE can do, but mentioned here for completeness) use different bridging inferences: **if the topic has been forbidden because of some aspect that is sensitive to the hearer, then**

- **Analogy:**

search for an analogue for the topic that contains an analogue of the sensitive part, and make an analogy. The fact that an analogy is being made can

be expressed explicitly by using phrases such as "It's just like...". The analogue, its parts, and its presentation, are, of course, also subject to the affect rule: there is no point in blindly substituting one sensitive topic for another.

- **Evidence:**

search for evidence for the new value of the sensitive aspect and, if applicable, evidence against the hearer's particular value of it, and say that. In this rule, *evidence for a value* means any fact, belief, or concept from which the hearer can infer the value. The fact that further inference is expected can be indicated by techniques such as ellipsis: "Your brother was involved in an accident and...". As before, the evidence itself and its presentation are subject to the affect rule.

4.4.2 Selectivity

A second topic inclusion strategy is *selectivity*. Sometimes, rather than evading unpleasant facts, the speaker just has to mention them directly and violate his goal not to offend the hearer. In such cases, though, he can be selective in his presentation: he can say appropriate aspects of topics that imply that the unpleasantness need not be considered important, or that the difference between what the hearer believes and what he is presenting is not too large. Most persuasive discussion is of this type.

This approach requires much more of the generator. In particular, it requires that the generator be able to search not only what it has been given to say, but also through the rest of memory, in order to find aspects of the topic that help achieve its pragmatic goals. The problem is how the generator must know where to search for such useful aspects.

Much work on topic collection for generation was done by McKeown (see, say, [McKeown 82, 83], [Paris & McKeown 87]). In her system, the program has goals to answer four different types of descriptive query to a database. Each goal has a set of strategies, called a schema, that contain instructions indicating which parts of the relevant database entry the generator should pick as answer. These strategies are, in fact, ossified plans, since they contain a fixed set of instructions in a fixed order. For example, one schema serves the goal

to *identify* (i.e., give a definition for) an object. This goal is activated in response to, for example, the request "what is a submarine?". When asked about a whisky-submarine, the *identify* schema directs the generator to say (produced by the TEXT generator, [McKeown 82]):

A whisky is an underwater submarine with a PROPULSION_TYPE of DIESEL and a FLAG of RDOR. A submarine is classified as a whisky if its CLASS is WHISKY. All whiskies in the ONR database have REMARKS of 0, FUEL_CAPACITY of 200, FUEL_TYPE of DIESEL, IRCS of 0, MAXIMUM_OPERATING_DEPTH of 700, NORMAL_OPERATING_DEPTH of 100 and MAXIMUM_SUBMERGED_SPEED of 15.

[McKeown 82, p 251]

In similar vein, PAULINE has three plans that indicate where in relation to the current topic it should search for additional sentence topics: the DESCRIBE, RELATE, and CONVINCE plans. The DESCRIBE plan is used to describe objects (McKeown's four schemas are specializations of this plan) and the RELATE plan to relate events and state-changes. The CONVINCE plan serves the goal of winning the hearer over to the speaker's opinion of the topic when they disagree.

In contrast to McKeown's schemas, the strategies contained in PAULINE's plans need not be applied in a fixed order. That is why the CONVINCE plan is called a plan and not a schema. When running the plan, PAULINE applies as many strategies as it has time for (depending on its rhetorical goal **RG:haste**), collects their results, and to them, in turn, applies an appropriate topic collection plan, until it runs out of time, or until no new candidates are found. Following that, the program is free to rearrange the candidate topics in order to achieve maximum effect, under guidance of other strategies based on the affect rule. The advantage of using a free-order plan instead of a schema is the additional flexibility it provides. PAULINE's plans simply make suggestions about what to say. The inclusion and rearrangement criteria are described in the next section; a more detailed overall description appears in chapter 7.

The CONVINCe plan serves the goal to support a specific affective opinion. In following it, the generator produces an argument in favor of its views. Each strategy of the plan helps to build up the argument by indicating *where in relation to the current topic the program can search for supporting evidence*.

Work in understanding and representing the structure of arguments, from which topic collection strategies can be derived, has been done by [Birnbaum, Flowers & McGuire 80], [Birnbaum 85] and by [Sycara-Kyranski 85]. In [Birnbaum 85] Birnbaum presents a scheme for representing arguments and for reasoning about the status (*established, in jeopardy*) of each of the argument propositions. He describes three rules by which propositions can support (or attack) other propositions in an argument: *appeal to authority, responsibility attribution*, and *justification*. Clearly, when the proposition the generator wants to establish is already represented in an argument graph of this type, with all the support and attack links made explicit, finding supporting topics is no problem. But since different hearers may find different lines of argument compelling, the generator must be able to assemble an argument supporting its opinion from scratch, taking into account the hearer's affects. What's more, it must be able to do so without relying on a hearer to make counterarguments that suggest new aspects as further topics. Thus PAULINE is conceived along the lines of a speechwriter, not a party in an argument.

When saying the argument, the generator should make clear what type of support or attack it is making; this is most easily achieved by beginning the sentence with an appropriate phrase — for example, "Well, Einstein said...", "Well, Pete's to blame for...", and "Well, don't you agree it's good that..." for the three types appeal to authority, responsibility attribution, and justification, respectively. The type of support or attack a proposition forms is determined by the relation its topic bears with the central topic of the argument: thus, for example, to find topics for responsibility attribution propositions, the generator must inspect the actors of concepts. In order to assemble an argument, then, the generator must test the speaker's and the hearer's affects for concepts with certain specific relations to the central topic.

Based on these considerations, as well as on analysis of various written arguments (taken

from communist newspapers, pro- and anti-labor strike leaflets, etc.) the CONVINCe plan contains the following strategies or suggestions for finding topics that support the speaker's version of a topic with a sensitive aspect (i.e., an aspect over which speaker and hearer have opposite affects):

The CONVINCe plan:

- Consider **worse examples** of the topic with the sensitive aspect — from the concept(s) immediately superior to the topic in the hierarchically organized memory network, compute the affects of other, similar instances, and collect those with affect equally BAD or worse
- Consider **good results** of the topic with the sensitive aspect — examine all the results and outcomes of the topic; if it is (part of) a goal, a plan, or a MOP (a stereotypical sequence of scenes; see [Schank 82]), examine the final outcomes too; collect those with GOOD affect
- Consider concepts with **good relations** to the topic with the sensitive aspect — compute the affects for the intergoal relations that the topic is part of (e.g., those goals the topic supports, opposes, is a side-effect of) and collect the GOOD ones
- Consider **good side-effects** of the topic with the sensitive aspect — examine all the side-effects of the topic (if it is a goal or a plan) as far as they are known in memory and collect the GOOD ones
- **Appeal to authority** — if any of the immediate aspects of the topic refer to people or organizations who share in, have, or support the sensitive aspect, and if the hearer's affect for these authorities is GOOD, collect them
- Simply **enhance or mitigate the topic** with the sensitive aspect — collect the topic and allow subsequent realization decisions to slant it appropriately

If the speaker wants to convince the hearer, he should look for aspects that the hearer agrees with and base his argument on that. However, when the hearer firmly holds the opposite opinion, this plan may not find anything on which they agree. In this case the speaker can simply follow his own sympathies, hoping that he will find something the hearer has not yet thought of that will cause him to change his opinion.

4.5 Form: Generating with Affect

4.5.1 Topic Organization

After it has collected a number of candidate topics, and before it says any of them, the generator must perform a number of *rhetorical planning* tasks to make its text elegant. Not only should it test the candidate topics for suitability as sentence topics, but it should examine the possibility of interpreting them, reordering them for maximum effect, and casting them into conjunctive phrases to make clear their individual roles in the text and their mutual relationships. Of course, the decisions made during rhetorical planning need not be based on affect alone. Other generators use a number of criteria to make these decisions; for example, *focus* in [McKeown 82]; *hearer knowledge* in [Cohen 78] and [Appelt 81]; *choosers* in [Mann 83a]; *hearer's opinion of speaker* in [Jameson 87]. But affect plays an important role as well, because much affect can be injected into the text during this stage of generation. This section contains a description of affect as a decision criterion.

Phrases

After reading the following paragraph, complete Martha's and Max's responses:

Martha and Max are little Pete's parents. Max and Pete are baseball fanatics, but Martha hates baseball. One day, Pete falls off his bicycle and is slightly hurt. Martha forbids him to play his baseball game that afternoon or to go

to the movies. Max, who wants his son to be tough, disagrees. Secretly, Pete sneaks out of the house and plays a splendid game, hitting five home runs, and then goes to see a movie. When he gets home, there is a fight between angry Martha and proud Max. The next day, Pete's grandmother calls, and asks both Martha and Max the same question: "So what has Pete done lately?". Max proudly says "He's been great; not only did he play baseball, but..."; and Martha angrily says "He's been bad; not only did he play baseball, but..."

Two appropriate responses are:

- (a) "Not only did he play baseball, but *he hit five home runs!*"
- (b) "Not only did he play baseball, but *he went to the movies afterwards!*"

Max's retort (a) implies that Pete's playing baseball was good, courageous, and tough. Martha's retort (b), in contrast, implies that it was bad and disobedient. *Each parent imputes an affect to Pete's playing by juxtaposing the sentence "he played baseball" with another sentence of suitable affect.*

Clearly, the "not only X but Y" sentence form is used to imply that X and Y carry the same affective value, and in fact that the value is to be strengthened due to their juxtaposition. In contrast, the sentences

- "Pete played the game and he hit five home runs"
- "When Pete played the game he hit five home runs"
- "Pete played the game. He hit five home runs"

carry no such cumulative affective import.

The "not only X but Y" form can be called an *enhancer*. More enhancing phrases:

- "Pete played the game; *also*, he hit five home runs"
- "Pete played the game; *in addition*, he hit five home runs"
- "Pete played the game; *what's more*, 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

- (i) 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, phrases with weakening effect are *mitigators*. 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:

- "*Although* John remembered the cat, he whipped the dog"
"John remembered the cat. *However*, he whipped the dog"
"John remembered the cat. *But still*, he whipped the dog"

Note that the first part, taken by itself, is neutral; it could just as well have been made BAD:

- "*Not only* did John remember the cat, he whipped the dog"

In a two-predicate mitigator, the sentence with the desired affect usually appears last. Two one-predicate mitigators are:

"Well, John whipped the dog, but..." (implying that other, as yet unknown, mitigating circumstances may exist)

"Oh well, so John whipped the dog" (implying that the action is unimportant)

A number of constraints must be met before two topics can be juxtaposed in an enhancer or mitigator phrase. Consider the following examples:

(c) "Not only did John feed the dog, but he saved the child's life"

(d) "Not only did John whip the dog, but he saved the child's life"

(e) "Not only did Pete play the game, but Sam hit five home runs"

Sentence (c) is fine: John's actions are GOOD in both parts. Sentence (d), however, is bizarre, because it is not clear which affect is to be strengthened. Therefore,

Constraint 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

In sentence (e), we hasten to find some relationship between Pete and Sam; we only accept the sentence if there is indeed some team-like conjunction of the two (and perhaps others) to carry the affect. Therefore,

Constraint 2:

Two-predicate enhancer and mitigator phrases can only be used when the topics in both parts have the same focus concept

Additional criteria for pairing topics in phrases are *topic similarity* and *derivation*. Using topic similarity, a BAD candidate is compared to the GOOD candidates and paired with the one with which it shares the most features, under a simple match that compares the representation types of corresponding aspects of the candidates. Thus, for example, (g) is a better match than (f):

(f) "Although John whipped the dog, he helped the old lady cross the road. He also fed the cat"

(g) "Although John whipped the dog, he fed the cat. He also helped the old lady cross the road"

because "dog" and "cat" are both animals. Alternatively, using the strategy of derivation, the planner may only cast two candidates into a phrase if they are members of the same subtree of topic collection; that is, if they were both collected when the convince plan was applied to their common parent topic. This corresponds to "sticking to the train of thought". In the shantytown example, PAULINE as a protester finds two topics it can use to enhance the destruction of the shantytown: the police arrest and the community's critical response. Since the former is closer to the destruction than the latter (its actor is also officials, its action is also a negation (albeit of someone's freedom rather than of something's construction), its object also directly involves students), the program casts it into the enhancer phrase, getting (h) rather than (i):

(h) NOT ONLY DID YALE HAVE OFFICIALS DESTROY THE SHANTYTOWN,
BUT THE POLICE ARRESTED 76 STUDENTS. THE COMMUNITY CRITICIZED
YALE'S ACTION.

(i) NOT ONLY DID YALE HAVE OFFICIALS DESTROY THE SHANTYTOWN, BUT
THE COMMUNITY CRITICIZED YALE'S ACTION. THE POLICE ARRESTED 76
STUDENTS.

Although PAULINE can use either strategy, the former tends to produce better text. This can be stated as a further constraint:

Constraint 3:

The predicates in two-predicate enhancer and mitigator phrases should match in as many aspects as possible, under a simple match of representation filler types

Chapter 7 contains a description of how the generator pairs up GOOD and BAD or GOOD and NEUTRAL topics, finds mitigating or enhancing phrases for each pair, and, after building up the goal to say the pair using the found phrase, sends the goal off to the generation routines to be said.

4.5.2 Sentence Inclusion

Adverbs

A number of adverbial stress words (such as “really”, “just”, and “only” for actions, and “very”, “extremely”, “slightly”, and “moderately” for states and adverbs) specifically function as enhancers or mitigators:

“Max *really* smashed Sam” — “Max *just* tapped Sam”

“John was *extremely* angry” — “John was *only* angry” (but not hysterical)

When these words are used to modify concepts that do not already carry affect, they seem strange, forcing the hearer to postulate affect; consider

“Mary merely looked at the book”

To remedy this, the constraint can be stated as

Constraint 4:

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

Thus affectively neutral words must not be modified by stress words ("really", "just"), but only by affect-laden adverbs ("hard", "narrowly"):

"Max hit Sam *hard*" — "Max *really* hit Sam"

"Sue *narrowly* won the race" — "Sue *just* won the race"

In addition, sentences such as

"Mike only killed Jim"

"Sue was merely ecstatic"

give rise to

Constraint 5:

Irreversible, extreme states and actions should not be mitigated

Verbs

Verbs play a very important role in the affective manipulation of text. Often, the verb used determines the content of the predicate, so that the selection of the verb is an important part of the sentence inclusion decision. For example, compare the affective difference between (a) and (b) for a Carter supporter:

(a) KENNEDY WON THE PRIMARY IN MICHIGAN

(b) KENNEDY BEAT CARTER IN THE PRIMARY IN MICHIGAN

Just as the mitigator phrase "X; however, Y" imputes to Y the opposite affect of X, the predicate forms of verbs can be viewed as having fields with characteristic affects. For example, for "beat",

"A beat B [in C]"

A, B, and C will be called *field fillers*. In the form for "beat", before anything is known about the field fillers, A carries the affect GOOD, B BAD, and C GOOD. These field affects can be used to impute affect, since the relative polarity of the fields remains unchanged: B must carry opposite affect to A and C. Seen from the perspective of verb choice, if the speaker likes Pete, he should not use "beat" (unless Bill and/or the race have previously been established as BAD):

"[That blighter] Bill beat [my hero] Pete in the [crooked] race"

Similarly, he should not use "lose" either, since in

"A lost [to B] [in C]"

field A carries BAD to fields B's and C's GOOD affects. However, he could use "win", which doesn't require a direct object, or "get", which avoids the losing altogether:

"Bill won the race"

"Carter got 20515 delegates"

and enables him to avoid placing his sympathy in a BAD field. This is what PAULINE does to produce (a) rather than (b) when it supports Carter in the example above. The incisive effect of this strategy (still supporting Carter) is even clearer when the sentence is embedded in a phrase:

(c) ALTHOUGH KENNEDY WON THE PRIMARY, CARTER IS AHEAD

(d) ALTHOUGH KENNEDY BEAT CARTER IN THE PRIMARY, CARTER IS AHEAD

Of course, field affects derive their existence from the semantics underlying the words such as "beat", "lose", and "win". Field affects are not *real* affects. Field affects are simply a concise way of representing the results of standard inferences the hearer is likely to make about the speaker's affects for the field fillers and of using them in generation. Winning is only GOOD, in neutral context, because of an assumed pre-existing set of goals; these are based on the concept WIN and not on the lexical item "win". Before using the field affects, the generator should check for the existence of such goals.

It is clear that, in order to select a verb with suitable affective predicate, the speaker must be able to inspect the predicate associated with each candidate verb. This can be

implemented in two ways: In one way, each verb furnishes its predicate(s) on demand, and the verb selection routines compare the field affect of each field in the predicate with the affect of the field filler, and select a verb whose predicate doesn't present a conflict. In the other way, the discrimination net that indexes the various verbs under the current representational item contains discriminations that test the various field fillers' affects and discriminate to an appropriate verb.

4.5.3 Sentence Organization

Affect has an effect on the organization of the parts of a sentence; specifically, on the choice of the subject and on the order of clauses.

In a typical sentence, almost any aspect of the sentence topic can be selected to be the sentence subject. Since it is a prominent position in the sentence, the subject must be chosen with care; random selection produces unconnected text:

(a) Jane and Susan went to see the new Spielberg movie last night.

The movie grossed \$10 million in its first two weeks.

They liked the movie a lot. — (a1)

The movie really enchanted them. — (a2)

It had been filmed in Morocco and California.

(b) Jane and Susan went to see the new Spielberg movie last night.

They were really looking forward to it.

The movie really enchanted them. — (b1)

They liked the movie a lot. — (b2)

Both Jane and Susan considered their money well spent.

In (a), the movie is established as the important topic; in (b), Jane and Susan are. Thus (a2) and (b2), with the movie and Jane and Susan respectively as sentence subjects, follow

more naturally than (a1) and (b1). When selecting an aspect of the input topic as sentence subject, the generator must be aware of these constraints.

Grosz [Grosz 77], Sidner [Sidner 79], and McKeown [McKeown 82], among others, have addressed this problem; the first two from the point of view of language analysis. Sidner and McKeown describe rules for choosing subjects in order to produce flowing, natural text. These rules are based on the notions *current focus* (the focus, usually, the subject, of the current sentence); *potential focus list* (the elements of a sentence that are candidates for the next sentence focus; in practise, most of its aspects); and *focus stack* (a most-recent-first list of the past sentence foci). McKeown's subject choice rules are, in order of preference,

1. select the new focus from the previous sentence's potential focus list
2. maintain the same focus as previous sentence
3. return to a topic of previous discussion from the focus stack
4. select the sentence topic with the most implicit links to the previous sentence's potential focus list

These rules are, of course, 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.

Using affect 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 — is another way of injecting affect into text. When the generator has the goal to convince the hearer or to make known its affects, it must use, in addition, the following rule to choose a sentence subject:

5. select the new focus from candidates with GOOD affect for sentences with GOOD affect and from candidates with BAD affect for sentences with BAD affect

4.5.4 Clause Content and Organization

Adverbs play a large role in communicating affect. Stress words were discussed above; in addition, other types of enhancer and mitigator adverbs (from the JUDGE domain) are:

intentionality: "intentionally", "purposely" — "accidentally"

degree: "hard" — "lightly" (hit)

number: "repeatedly" — "once" (stabbed)

During the realization of a sentence, the speaker must find the adverbs he can legitimately say (one cannot, for example, misrepresent the contents of the topic to say "lightly" when the aspect DEGREE is HARD), and choose some, usually at most two (when affective adverbs are overused the effect is unnatural; PAULINE has been limited, arbitrarily, to two per sentence):

- (a) MIKE JUST HIT JIM ONCE.
- (b) JIM COULD FORESEE THAT HE WOULD HURT MIKE IF HE REALLY
SLAPPED HIM. HE SLAPPED HIM REPEATEDLY AND HURT HIM.

Within a clause, the speaker has to decide which aspects of the topic to say and how to order them. For example, when making a noun group, he must select the head noun and then decide whether to describe it in full, only give unsaid information, or give an abbreviated version. He then has to select and order the modifiers, both pre- and post-nominal (and some modifiers can appear in both positions), before he can construct a form from which the eventual noun group will be built. These decisions are determined by the speaker's rhetorical goals of style (as described in chapter 5, the level of formality plays a role), and of opinion (affect plays a role too). For example, when the speaker has affect for the object, his opinion may be expressed by the head noun (a), by an adjective (b), or by a post-nominal modifier (c):

- (a) THAT RAT, MIKE
- (b) THAT WONDERFUL MIKE
- (c) MIKE, THE GENTLEMAN

4.5.5 Word Choice

The specific words used contribute greatly to the affective content of the text.

Verbs

In addition to determining the form of the predicate, as discussed above, verbs themselves often carry some affect. Often an action can be described by a number of verbs; for example, some enhancing/mitigating verbs are:

"Jane *slammed/tapped* me on the head with a crowbar"

"Mike *wolfed down/nibbled* his supper"

A nice result of constraint 4 is that it helps to organize related words such as "hit", "smash", and "tap" in the lexicon. The sensible way to use these three words is to access "hit" from the representation HIT and then perform a (series of) affective discrimination(s) until the appropriate word "smash" or "tap" is found. (This idea was discussed in [Goldman 75], though his discriminations depend on features of the input concept rather than on the generator's affective goals: INGEST with a liquid OBJECT gives "drink", with solid OBJECT "eat", and with gaseous OBJECT "breathe".) The problem is to determine *which words are affectively neutral and therefore good starting points for the discrimination*; while this may be immediately clear for "hit" and its variants, it is less so for "punch" or "slap". Chris Owens [personal communication 85] suggests that the adverbial stress words can test speakers' intuitions about the affective neutrality of a given word. Of the following sentences, *said without stressed intonation*,

- (a) "Pete really hit Joe"
- (b) "Pete really hit Joe hard"
- (c) "Pete hit Joe really hard"
- (d) "Pete really punched Joe"
- (e) "Pete just slapped Joe"

(a) feels awkward and therefore fails the test, so "hit" is neutral. Both (b) and (c) pass it ((c) better than (b), which indicates just where the affect lies). To me, (d) also feels awkward, but (e) seems fine; thus to me "punch" is affectively neutral but "slap" is not. PAULINE's lexicon is organized on these lines; for example, from the action HIT it accesses "hit" and then discriminates to "tap" as a mitigator and "smash" as an enhancer:

(a) JIM JUST TAPPED THAT JERK MIKE ONCE

(b) JIM PURPOSELY SMASHED MIKE AND KNOCKED HIM DOWN

Of course, there is more to sentence affect than this. Sometimes a sentence as a whole carries affect, though each individual word is neutral. In such cases, stress words can be used too. For example,

"The big man hit his aged mother in the mouth"

is certainly not neutral. The sentence creates a sense of incredulity. Applying the stress-word test again, the sentence

"The big man really hit his aged mother in the mouth"

accents exactly this feeling, rather than the action "hit". Much work remains to be done in this regard.

In the shantytown examples, PAULINE uses "tear down" and "destroy" as enhancers and "remove" as a mitigator for the action of disassembling the shanties. For communication of a request, the lexicon contains "order" and "command" as enhancers and "request" and "ask" as mitigators; the neutral word is "tell". Compare the following extracts:

(c) YALE ORDERED THE STUDENTS TO BUILD THE SHANTYTOWN ELSEWHERE.
LATER, AT 5:30 AM ON APRIL 14, OFFICIALS DESTROYED THE
SHANTYTOWN.

(d) YALE REQUESTED THAT THE STUDENTS BUILD THE SHANTYTOWN
ELSEWHERE, BUT THEY REFUSED TO LEAVE. THE SHANTYTOWN WAS REMOVED

FROM BEINECKE PLAZA BY OFFICIALS ONE MORNING.

In the Carter-Kennedy example, PAULINE uses "cream", "trounce", and "triumph" as versions of "beat" (all these are enhancers; further discrimination depends on the desired level of formality and floridity, as discussed in chapter 5). Variants for the concept "narrow lead" are "narrow", "reduce", "decrease", or "diminish" the lead, as well as "suffer setback" and "lose ground":

(e) KENNEDY DIMINISHED CARTER'S LEAD BY GETTING ALL OF 21850 VOTES IN THE PRIMARY IN MICHIGAN. IN A SIMILAR CASE, CARTER DECREASED UDALL'S LEAD IN A PRIMARY IN 1976, AND HE EASILY TROUNCED UDALL TO BE NOMINATED BY 2600 DELEGATES. I AM REAL GLAD THAT PRESENTLY KENNEDY IS CLOSER TO GETTING THE NOMINATION THAN BEFORE.

..

Adjectives

Of course, adjectives can also carry affect:

enhancer: "The *wimpy* boy rode home" (compared to "the small boy")

mitigator: "The actress was convincing; she is *adequate*"

(compared to "she is fine")

Depending on the desired level of partiality, formality, and floridity (as described below and in chapter 5), PAULINE can express its opinion by including appropriate adjectives, such as "concerned", "wonderful", "nice", and "nasty".

Nouns

Nouns can also express the speaker's opinions. For example, saying "terrorist" to an IRA soldier may get you shot; saying "freedom fighter" certainly will not! PAULINE can express its opinion of an actor when building a noun group by selecting a mitigator or enhancer noun; for example, depending on affect and on the desired level of partiality and formality, it can say:

- (a) THAT JERK, JIM, ...
- (b) THAT GENTLEMAN, JIM, ...

4.6 When and Where to be Partial

When the speaker's sympathies differ from the hearer's over the topic, the speaker faces a potential problem. If he says the topic he may alienate the hearer, something he may not want to do. He can decide not to say the topic at all. If he does, however, he must decide how explicit to be: he can merely refer to sensitive aspects obliquely, he can mitigate them somehow, or he can say them straight out.

4.6.1 Timidity: The Inclusion of Sensitive Topics

There are at least three points in the generation process when the speaker can decide to ignore sensitive topics:

- **collection of new candidate sentence topics:** Whatever its origin — a topic collection plan, input, a reminding — a sensitive topic or aspect of a topic can simply be excluded from further topic collection. As a result, none of its results, side-effects, or any other aspects will appear in the subsequent text
- **interpretation of candidate topics:** a sensitive topic can be subsumed by a high-level interpretation of it, if one can be found, as described in chapter 3
- **construction of phrases that contain topics:** a sensitive topic can be juxtaposed with others in a phrase that mitigates its affect, if an appropriate one can be found

The speaker's willingness to include affectively sensitive topics depends not only on his and the hearer's sympathies, but also on aspects of the conversation setting, such as the amount of time available. The less time and the less stylistic flexibility, the less opportunity the

speaker has to present the topics in a way that will achieve his goals. Under appropriate circumstances, the speaker may simply decide not to include sensitive topics at all.

These considerations give rise to the following general rule: the speaker should only include sensitive topics if he is not overly concerned with the hearer's reaction — that is, if he is socially dominant, if the hearer is a stranger, or if the hearer is distant from him anyway. This rule can be made more precise by stating it in terms of the pragmatic features given in chapter 2 — that is, by giving the rules PAULINE uses to activate the rhetorical goal **RG:timidity** that controls the inclusion of sensitive topics. The goal takes a value from the list *timid*, *normal*, *reckless* according to the following rules:

1. set **RG:timidity** to *timid* if speaker's and hearer's affects for the topic do not agree, and if: **desired effect on hearer's emotion toward speaker** is marked *like*; or if **speaker-hearer relative social status** is marked *subordinate*, that is, the speaker is socially subordinate to the hearer
2. set **RG:timidity** to *reckless* if: **speaker's interest level** is marked *high*, and the **desired effect on interpersonal distance** is marked *distant*, that is, that the speaker is very interested the topic and doesn't care much about the hearer's opinion of him; or if **speaker-hearer relative social status** is marked *dominant*, since usually a dominant speaker expresses his opinions more freely than a subordinate one; or if **desire to involve hearer** is marked *involve*, since by being more reckless he may goad the hearer into argument or discussion
3. otherwise, set **RG:timidity** to *normal*

When activated with the value *timid*, this goal causes PAULINE not to incorporate sensitive aspects encountered during topic collection, interpretation, and phrase construction. With the value *reckless*, the opposite occurs; with the intermediate value, PAULINE includes sensitive topics only if certain other rhetorical goals (such as **RG:haste**) permit.

4.6.2 The Degree of Partiality

How does the degree of partiality in the text affect the hearer? Clearly, this depends on the hearer's own opinions. If the hearer is a protester himself, he will react very differently to the following two sentences:

(a) AS A REMINDER TO YALE UNIVERSITY TO DIVEST FROM COMPANIES WITH BUSINESS IN SOUTH AFRICA, A LARGE NUMBER OF CONCERNED STUDENTS BUILT A SHANTYTOWN, WINNIE MANDELA CITY, ON BEINECKE PLAZA ONE DAY.

(b) A FEW SHIFTLESS STUDENTS TOOK OVER BEINECKE PLAZA ONE DAY AND BUILT A SHANTYTOWN, WINNIE MANDELA CITY, BECAUSE THEY WANTED TO FORCE YALE UNIVERSITY TO DIVEST FROM COMPANIES WITH BUSINESS IN SOUTH AFRICA.

Obviously, to him, the latter is more contentious. Statements like (b) cause fights. In order to know how explicitly partial to be, the speaker must know a lot about the hearer, the topic, and the conversational circumstances. In general, if he wants to remain friendly with the hearer, his conduct must be governed by the affect rule stated above; thus, when the speaker's opinion agrees with the hearer's, expressing it will tend to make them closer; when it disagrees, expressing it may cause fights. These considerations can be stated concisely as the following rules (defined in terms of the pragmatic aspects given in chapter 2) which PAULINE uses to assign its rhetorical goal **RG:partiality** a value from the list *explicit*, *implicit*, *impartial*. This rhetorical goal controls how explicitly the generator states its opinions, as described above:

1. set **RG:partiality** to *explicit* if the speaker's and hearer's **affects for the topic agree** and: **desired effect on hearer's emotion toward speaker** is marked *like*; or if **desired effect on interpersonal distance** is marked *close*; or if **desired effect**

on hearer's emotional state is marked *calm*, that is, the speaker doesn't want to anger the hearer; or if **speaker-hearer relative social status** is marked *equal* or *subordinate*, that is, if the speaker does not want to imply social domination over the hearer; or if **atmosphere (tone)** is marked *informal*, in order to relax the tone of the conversation

2. set **RG:partiality** to *implicit* if the speaker's and hearer's **affects for the topic agree** and: **desired effect on interpersonal distance** is marked *distant*, since being lukewarm about their agreement with the hearer separates them; or if **speaker-hearer relative social status** is marked *dominant*, for the same reason; or if **desire to involve hearer** is marked *leave*, that is, if the speaker does not want get the hearer involved in the conversation; or if **speaker's desire to emphasize** is marked *mention*, in contrast to emphasis
3. otherwise, set **RG:partiality** to *impartial* if their **affects agree**, or if their **affects disagree** and **hearer's knowledge level** is marked *expert* and **speaker's knowledge level** is marked *student* or *novice*, and **desired effect on hearer's emotion toward speaker** is marked *respect* or *like*, since when the speaker cares about an expert hearer's opinion of him, he will not want to exhibit his partiality and lack of knowledge
4. set **RG:partiality** to *explicit* if the speaker's and hearer's **affects for the topic disagree** and: **desired effect on hearer's emotional state** is marked *anger*; or if **desired effect on hearer's emotion toward speaker** is marked *dislike*
5. otherwise, set **RG:partiality** to *implicit* if their **affects disagree** and: **desire to involve hearer** is marked *involve*; or if **speaker-hearer relative social status** is marked *subordinate* (that is, the hearer is subordinate)

4.7 Conclusion

This chapter describes some techniques a speaker can use to incorporate his opinions in his text. To do so, he must be able to determine his opinion about any piece of the topic; this can be achieved by following rules governing the combination and propagation of his affects along the relationships among representation elements. The speaker can emphasize or evade topics he likes or dislikes by being selective in his choice of topic, by juxtaposing topics of appropriate affect within affect-imputing phrases, by selecting verbs with appropriate predicate forms, by including appropriate stress words, adverbs, and adjectives, and by choosing verbs and nouns with appropriate affect. All these decisions are based upon the affect rule.

In summary, again, consider the derivation of a paragraph from the JUDGE example, when PAULINE is defending Mike:

AFTER THAT, JIM DID NOT EXPECT	topic: NEUTRAL
THAT [THAT JERK] MIKE WAS GOING	enhancer BAD noun
TO HURT HIM ANY LONGER;	
[ALSO], JIM COULD FORESEE THAT	enhancer phrase
HE WOULD INJURE MIKE IF HE	topic: GOOD for Mike
[PURPOSELY] STRUCK HIM.	enhancer intent
HE [REALLY] [SMASHED] HIM.	enhancer stress, verb
THE RESULT WAS THAT HE INJURED HIM.	topic: GOOD for Mike

as well as when PAULINE supports Kennedy:

IN THE PRIMARY IN MICHIGAN, KENNEDY [DIMINISHED] CARTER'S LEAD BY GETTING [ALL OF] 21850 VOTES. IN A SIMILAR CASE, CARTER DECREASED UDALL'S LEAD IN A PRIMARY IN 1976, AND HE [EASILY] [TROUNCED] UDALL TO BE NOMINATED BY 2600 DELEGATES. [I AM REAL GLAD THAT] PRESENTLY KENNEDY IS CLOSER TO GETTING THE NOMINATION THAN BEFORE.	topic: GOOD for Kennedy enhancer verb enhancer topic: GOOD reminding enhancer stress word enhancer verb explicit opinion topic: GOOD for Kennedy
---	---

Also compare two versions of the shantytown episode (both formal):

FOR PROTESTERS	FOR UNIVERSITY	
IN EARLY APRIL, [AS A REMINDER TO] YALE UNIVERSITY TO DIVEST FROM COMPANIES DOING BUSINESS IN SOUTH AFRICA, A [LARGE NUMBER OF] [CONCERNED]	IN EARLY APRIL, [A SMALL NUMBER OF] STUDENTS [TOOK OVER] BEINECKE PLAZA AND CONSTRUCTED A SHANTYTOWN, WINNIE MANDELA CITY, [IN ORDER TO FORCE] YALE UNIVERSITY TO DIVEST FROM COMPANIES DOING BUSINESS IN SOUTH AFRICA.	mitigator enhancer/mitigator interpretation interpretation
STUDENTS CONSTRUCTED A SHANTYTOWN, WINNIE MANDELA CITY, ON BEINECKE PLAZA.		
THE LOCAL COMMUNITY EXPRESSED SUPPORT FOR THE STUDENTS' ACTION.	YALE REQUESTED THAT THE STUDENTS MOVE THE SHANTYTOWN, BUT THEY REFUSED TO LEAVE. THE UNIVERSITY INTENDED TO BE REASONABLE. YALE GAVE IT PERMISSION TO EXIST UNTIL THE MEETING OF THE YALE CORPORATION, BUT EVEN AFTER THAT THE STUDENTS STILL REFUSED TO MOVE.	topic: support topic: request explicit opinion topic: permission topic: refusal

PROTESTER, continued	UNIVERSITY, continued	
LATER, AT 5:30 AM ON APRIL 14, THE SHANTYTOWN WAS DISASSEMBLED BY OFFICIALS; ALSO, AT THAT TIME, THE POLICE ARRESTED 76 STUDENTS. THE STUDENTS REQUESTED THAT YALE GIVE THEM PERMISSION TO REASSEMBLE IT ON BEINECKE PLAZA AND AT THE SAME TIME SEVERAL LOCAL POLITICIANS AND FACULTY MEMBERS EXPRESSED CRITICISM OF YALE'S ACTIONS. FINALLY, YALE PERMITTED THEM TO RECONSTRUCT THE SHANTYTOWN.	AT 5:30 AM ON APRIL 14, OFFICIALS [HAD TO] DISASSEMBLE THE SHANTYTOWN. FINALLY, YALE [BEING CONCILIATORY TOWARD THE STUDENTS], [NOT ONLY] PERMITTED THEM TO RECONSTRUCT THE SHANTYTOWN, [BUT ALSO] ANNOUNCED THAT A COMMISSION WOULD GO TO SOUTH AFRICA IN JULY TO EXAMINE THE SYSTEM OF APARTHEID.	interpretation topic: arrest topic: request topic: support interpretation enhancer topic: commission

Chapter 5

Creating Style

Abstract

This chapter addresses the question “how do we say the same thing in different ways in order to communicate non-literal, pragmatic information?”. Clearly, the style of the realized text — formal, hasty, forceful — plays a large role in its efficiency in achieving the speaker’s goals. Any generator worth its salt must be able to produce text of different styles; in order to do so, it requires rules that determine the types of options it must choose at decision points in the topic organization and text realization processes. This chapter takes an algorithmic approach to the creation of style in language.

5.1 The Nature of Style

Text style contains pragmatic information: it may be stuffy, slangy, prissy, etc. In order to produce pragmatic-based, goal-directed language, generators have to be able to manipulate style: therefore, they require an understanding of what it is, what effects various styles have on the hearer, and what information various styles convey. Since classifying all possible styles of text is an impossible task, a number of basic styles can be identified from which

more sophisticated and idiosyncratic styles can be built, as discussed in chapter 2. Typically, handbooks of good writing describe styles in terms of the characteristics of complete paragraphs of text, which is not very useful for a practical, generator-oriented approach. Instead, a functional approach is to describe styles in terms of the decisions a generator has to make: decisions such as sentence content, clause order and content, and word selection.

During the construction of PAULINE, the following fact was found empirically: when you vary the decisions you make during generation, certain types of decisions group together and form stylistically coherent text, and other types, when grouped, produce text that is incoherent or odd. The coherent groupings conform to traditional stylistic concepts. This chapter illustrates a few such coherent groupings, and hence develops functional definitions for the text styles **formality**, **haste**, and **force**.

5.2 Formality

The level of formality of text is probably the most obvious stylistic aspect, since it is a criterion that plays a role along the whole range of generator decisions — from the initial topic selections and organization down to the final word selection. All language users have rules for making their text more or less formal; some people, such as politicians, have a large number of such rules and hence produce good examples of formal language. The best way to illustrate these rules is to dissect a piece of text:

Yesterday, December 7, 1941 — a date which will live in infamy — the United States of America was suddenly and deliberately attacked by naval and air forces of the Empire of Japan.

The United States was at peace with that nation and, at the solicitation of Japan, was still in conversation with its Government and its Emperor looking forward to the maintenance of peace in the Pacific.

Indeed, one hour after Japanese air squadrons had commenced bombing Oahu, the Japanese Ambassador to the United States and his colleague delivered to the Secretary of State a formal reply to a recent American message. While this reply stated that it seemed useless to continue the existing diplomatic negotiations, it contained no threat or hint of war or armed attack.

It will be recorded that the distance of Hawaii from Japan makes it obvious that the attack was deliberately planned many days or even weeks ago. During the intervening time, the Japanese Government has deliberately sought to deceive the United States by false statements and expressions of hope for continued peace.

[“We Will Gain the Inevitable Triumph — So Help Us God”, war address by U.S. President F.D. Roosevelt to joint session of Congress of the United States, December 8, 1941.]

What characteristics make this address formal? Certainly, one factor is the use of formal verbs and nouns instead of more common ones, such as “solicitation” instead of “request”. Another factor is the use of full names and titles instead of their common abbreviations. Replacing words and phrases in the address by less formal equivalents ((a) below), and using the common names for entities (b), we get:

Yesterday, December 7, 1941 — a date which will live in infamy — the US_(b) was suddenly and deliberately attacked by naval and air forces of Japan_(b).

The US_(b) was at peace with that nation and, at the request_(a) of Japan, was still talking to_(a) its Government and its Emperor looking forward to the maintenance of peace in the Pacific.

Indeed, one hour after Japanese air squadrons had started_(a) bombing Oahu, the Japanese Ambassador here_(a) and his colleague gave_(a) the Secretary of State a formal reply to a recent American message. While this reply said_(a) that it seemed useless to continue the existing diplomatic negotiations, there was_(a) no threat or hint of war or armed attack.

It will be recorded that the distance of Hawaii from Japan makes it obvious that the attack was deliberately planned many days or even weeks ago. In_(a) the intervening time, the Japanese Government has deliberately tried_(a) to cheat_(a) the United States by false statements and expressions of hope for continued peace.

However, the result is definitely not informal. The sentences still seem long and involved. In order to simplify them, we (a) remove conjunctions and multi-predicate phrases, and (b) remove adverbial clauses, or place them toward the ends of sentences (note the difference this change makes to the very first sentence):

The US was suddenly and deliberately attacked by naval and air forces of Japan yesterday, December 7, 1941_(b). This date will live in infamy_(a,b).

The US was at peace with that nation. [and_(a)] At the request of Japan, the US was still talking to its Government and its Emperor looking forward to the maintenance of peace in the Pacific.

Indeed, one hour after Japanese air squadrons had started bombing Oahu, the Japanese Ambassador here and his colleague gave the Secretary of State a formal reply to a recent American message. [While_(a)] This reply said that it seemed useless to continue the existing diplomatic negotiations. [, there_(a)] There was no threat or hint of war or armed attack.

It will be recorded that the distance of Hawaii from Japan makes it obvious that the attack was deliberately planned many days or even weeks ago. The Japanese Government has deliberately tried to cheat the United States by false statements and expressions of hope for continued peace in the intervening time_(b).

Though not a formal address any longer, the text is not yet informal; in fact, it seems odd. For example, phrases such as "looking forward to the maintenance of peace" and "it will be recorded" do not blend with phrases such as "deliberately tried to cheat". In order to improve this, we (a) eliminate the use of passive voice, and (b) refer to the involved parties — speaker, hearer, and others — directly:

We_(b) were suddenly and deliberately attacked by naval and air forces of Japan yesterday, December 7, 1941. This date will live in infamy.

We_(b) were at peace with that nation. At the request of Japan, we_(b) were still talking to their_(a) Government and its Emperor. We were_(b) looking forward to the maintenance of peace in the Pacific.

Indeed, one hour after Japanese air squadrons had started bombing Oahu, their_(b) Ambassador and his colleague gave our_(b) Secretary of State a formal reply to a recent message. This reply said that they thought it was_(a,b) useless to continue the existing diplomatic negotiations. But they made_(a,b) no threat or hint of war or armed attack.

Note_(a,b) that the distance of Hawaii from Japan makes it obvious that they deliberately planned_(a) the attack many days or even weeks ago. The Japanese Government has deliberately tried to cheat us_(b) by false statements and expressions of hope for continued peace in the intervening time.

Now some phrases sound flowery and out of place. To simplify, some nominalized verbs are converted back to verbs (a); noun groups are simplified by dropping redundant adjectives and nouns (b); pronominalization is increased (c):

We were suddenly and deliberately attacked by naval and air forces of Japan yesterday, December 7, 1941. This date will live in infamy.

We were at peace with them_(c). At Japan's request_(b), we were still talking to their Government and its_(b) Emperor. We were looking forward to maintaining_(a) peace in the Pacific.

Indeed, one hour after Japanese air squadrons had started bombing Oahu, their Ambassador and his colleague gave our Secretary of State a formal reply to a recent message. This reply said that they_(c) thought it was useless to continue negotiating_(a). But they_(c) made no hint of armed attack_(b).

Note that the distance of Hawaii from Japan makes it obvious that they deliberately planned the attack many days or even weeks ago. The Japanese

Government has deliberately tried to cheat us by false statements and_(b) expressions of hope for continued peace in the intervening time.

Finally, a few finishing touches: simplified tenses (a); colloquial phrases (b); elision of redundant words where grammatical (c):

We were suddenly and deliberately attacked by naval and air forces of Japan yesterday, December 7, 1941. We'll never forget this date_(b,c).

We were at peace with them. At Japan's request, we were still talking to their Government and Emperor. We were looking forward to having_(b) peace in the Pacific.

[Indeed_(c)] one hour after Japanese air squadrons had_(a) started bombing Oahu, their Ambassador [and his colleague_(c)] gave our Secretary of State a formal reply to a recent message. [This reply said that_(c)] they thought it was useless to continue negotiating. But they didn't_(b) hint at armed attack.

[Note that_(c)] The distance of Hawaii from Japan makes it obvious that they deliberately planned the attack a while_(b,c) ago. The Japanese Government has_(a) deliberately tried to cheat us by false expressions of hope for continued_(b) peace in the intervening_(b) time.

Though not colloquial yet, more changes of the same kind (such as using slang and idioms, and inserting explicit opinions) can change the address even further. As shown, the changes have produced informal text, something Roosevelt might have said to his family but not to Congress:

We were suddenly and deliberately attacked by naval and air forces of Japan yesterday, December 7, 1941. We'll never forget this date.

We were at peace with them. At Japan's request, we were still talking to their Government and Emperor. We were looking forward to having peace in the Pacific.

One hour after Japanese air squadrons started bombing Oahu, their Ambassador gave our Secretary of State a formal reply to a recent message. They thought it was useless to continue negotiating. But they didn't hint at armed attack.

The distance of Hawaii from Japan makes it obvious that they deliberately planned the attack a while ago. The Japanese Government deliberately tried to cheat us by pretending to hope for peace in the mean time.

As the sequence of transformations shows, the intermediate texts seem odd: they contain phrases and words that do not blend together well. Thus, all the steps are necessary to create stylistically coherent text. The rules underlying the transformations operate together, as a unit, to manipulate the level of formality.

The transformation rules were gathered from the analysis of a number of texts, ranging from politicians' speeches and writings to discussions with friends. In order to make text more formal, a speaker or generator must examine the options at each of the following decision points and apply the following strategies:

- **topic inclusion:** select options that contain causal, temporal, or other relations to other sentence topics, so as to make long sentences.
- **topic organization:** select options make long, complex sentences, by subordinating them in relative clauses; by conjoining two or more sentence topics; by juxtaposing them into relations and multi-predicate enhancer and mitigator phrases
- **sentence organization:** make sentence seem weighty by including a number of adverbial clauses; by placing these clauses toward the beginnings of sentences rather than at the end; by building parallel clauses within sentences; by using passive voice; by using more "complex" tenses such as the perfect tenses; by avoiding ellipsis, even though it may be grammatical, in such sentences as "Joe got more than Pete [did]", "When I was 20 years old, I got married for the first time"

- **clause organization:** make weighty, formal clauses, by including many adjectives and adjectival clauses in noun groups; by doubling nouns in noun groups ("Government and Emperor", "statements and expressions"); by including many adverbs and stress words in predicates; by using long, formal phrases; by nominalizing verbs and adverbs as in "their flight circled the tree" instead of "they flew round the tree"; by pronominalizing where possible; by not referring directly to the interlocutors or the setting
- **phrase/word choice:** select formal phrases and words; avoid doubtful grammar, popular idioms, slang, and contractions (for example, avoid slang by saying "man" rather than "guy" and contractions with "cannot" rather than "can't")

In contrast, the generator can make text less formal by

- **topic organization:** make simple sentences, by not conjoining two or more sentence topics; by not juxtaposing them into relations or affect-imputing phrases; and by not subordinating them in relative clauses
- **sentence organization:** make simple sentences, by including at most one adverbial clause per sentence, placed toward the end of the predicate; by using active voice; by avoiding the perfect and other "complex" tenses; by eliding words and clauses where this may grammatically be done
- **clause organization:** make simple clauses, by selecting at most one adjective in noun groups; by using short, simple phrases; by pronominalizing where possible; by using verbs and adverbs instead of their nominal forms; by referring to the interlocutors and the setting directly
- **word choice:** use informal phrases and words; by selecting only simple, common words; by using popular idioms, slang, and contractions wherever possible

Knowing how to be formal is not enough. The generator must also know when formal language is appropriate. Usually, interlocutors establish some level of formality at the

outset of their interaction: even the range of greetings "good morning", "hello", and "hi" helps to indicate a desired level. Once established, the level of formality doesn't normally change very quickly, and since formality is not precisely measurable, it is most apparent only when the level is suddenly changed or is inappropriate. In order to see the pragmatic effects of using formal language, then, the important question is: *what does the speaker achieve by altering the level of formality?*

First, if you become less formal, you signal a perceived or desired decrease in the interpersonal distance between yourself and the hearer. In any relationship, the people involved maintain a certain distance (say, on a range from intimate to aloof). Mirroring this, in conversations, the speaker and the hearer usually adopt some level of formality that, to them, accurately reflects the distance. Which interpersonal distance corresponds to which level of formality depends, of course, on social convention and on the interlocutors and their relationship; for example, colloquial or informal language is normally used to discuss relatively intimate topics. (The odd feeling produced by exceptions — say, when psychiatrists and patients converse about intimate topics in formal language — confirms this rule.) Usually, a lower degree of formality permits the selection of more intimate topics and the use of more personal phrases and words (and, of course, more slang, more flexible interruption behavior (see [Sacks, Schegloff & Jefferson 74]), etc.). Conversely, greater formality indicates that you feel, or wish to feel, more distant than the conversation had been implying, perhaps after the hearer had offended you, or when you had become uncomfortable with the topic. See [Brown & Levinson 78] on the use of language in formal situations, and [Kuno 73], [Harada 76], [Gasser & Dyer 86] on Japanese deictic honorifics.

Second, if you alter the level of textual formality, you may perturb the tone or atmosphere of the conversation. Since the level of textual formality roughly parallels that of the conversational atmosphere, a serious conversation such as a speech at a burial or a talk at a conference requires more formality than an everyday conversation such as a report to the family of the day's events. (This remains true when a new topic calls for a change in tone, even if the hearers and the setting do not change. To help signal such a change, speakers often use "important" voices, clear their throats, or speak more softly and slowly.) In such cases, an inappropriate level of formality can affect the hearer's emotion toward

you: for example, if you are too informal, you may seem cheeky or irreverent; if you are too distant, you may seem snooty or cold. As mentioned in chapter 2, a large amount of work by sociologists, anthropologists, and psycholinguists describes the characteristics of various settings and the appropriate levels of formality in various cultures (see, for example, [Irvine 79] and [Atkinson 82] on formal events; [Goody 78] and [R. Lakoff 77] on politeness).

Based on these consideration, speakers seem to use the following rules to determine an appropriate level of formality (PAULINE's rhetorical goal **RG:formality** takes a value from the range *highfalutin*, *normal*, *colloquial*):

1. set **RG:formality** to

- *colloquial* when the **depth of acquaintance** is marked *friends*, or when the **relative social status** is marked *equals* in an **atmosphere (tone)** marked *informal*
- *normal* when the **depth of acquaintance** is marked *acquaintances*
- *highfalutin* when the **depth of acquaintance** is marked *strangers*

2. then, reset **RG:formality** one step toward *colloquial* if: **desired effect on interpersonal distance** is marked *close*, that is, if the speaker wants the hearer to feel closer to him; or if **tone** is marked *informal*, that is, if the conversation occurs in a relaxed, friendly atmosphere

3. or reset **RG:formality** one step toward *highfalutin* if: **desired effect on interpersonal distance** is marked *distant*, that is, if the speaker wants to increase the emotional distance between himself and the hearer; or if **tone** is marked *formal*, that is, if the speaker wants to establish a serious tone for the conversation, or if he is making a speech at a formal occasion

4. and invert the value of **RG:formality** if: **desired effect on hearer's emotion toward speaker** is marked *dislike*, since inappropriate formality is often taken as an insult; or if **desired effect on hearer's emotional state** is marked *angry*. (The

contrapositive of these two rules provides the default rule: to make the hearer like you, select an appropriate level of formality.)

5.3 Haste

Haste refers to the amount of time the speaker allows himself to generate language. The less time available, the more pressure on the speaker, the less effort he can spend in making his text appropriate and striking. In order to decrease the time he needs, the speaker can minimize processing of non-essential tasks in the generation process or even ignore them altogether. Thus the rhetorical goal of haste affects the decision points that occur where the speaker can take short-cuts.

How can one determine where such short-cut points are? One way is to examine what mistakes people make when they speak under pressure. Some typical excuses are:

- *"I ran out of things to say — I said it all immediately, and then all I could do was repeat myself"* — no **additional topic collection** is performed before the main topics are all realized. In this respect, haste controls the extent to which additional topics are gathered under guidance of topic collection plans such as the CONVINCE plan described in chapter 4.
- *"I talked about something I shouldn't have"* — unhappy **topic choice**, alias foot-in-the-mouth disease. Since it is unlikely that the speaker will know beforehand the hearer's opinion on all aspects of the topic, he will have to determine the likely effect of saying it by guessing the hearer's relevant sympathies. This process may take some time, especially if the topic is complex. Similarly, the stylistic appropriateness of a sentence topic may sometimes be difficult to determine. Since under pressure people perform this task badly, their sense of haste must help determine the amount of affective and stylistic checking and censorship performed.
- *"I used the wrong word or phrase"* — inappropriate **word and phrase choice** is a common problem: well-planned text is characterized by apt and imaginative

expressions. As for topics, checking the affective implications of words and phrases can take time.

- *"I got all muddled in my sentence and couldn't say what I meant"* — when the sentence topic is complex and related to other topics in a large number of ways, badly done or neglected **sentence organization** can cause uncompletable sentences, elucubration, unavoidable repetitions, etc. The goal to be hasty can be used to control how much sentence planning is done.
- *"I just said whatever came into my head"* — inadequate **topic organization** takes many forms: as described in chapter 7, sentence topics must be ordered to form coherent paragraphs; they must be linked by introductory and other phrases that express the relationships between adjacent sentence topics; they can be juxtaposed into multi-predicate phrases (of affect, as described in chapter 4, or of spatial, temporal, and other relations, as described in chapter 6); also, as argued in chapter 3, generators should be able to activate inferences that make interpretations of topics and check the interpretations for suitability. Before performing any of these tasks, the generator should check whether it has the time to do so.

Errors of these types are common in everyday speech. For example, consider the following transcript of a conversation between three young women on the question *the current trend towards increased employment of women must ultimately conflict with the child care needs of society*. Undoubtedly the three interlocutors experienced very high pressure to say something -- not only were they strangers to each other, but the conversation was being recorded as an experiment, and they had been instructed to arrive at some common conclusion:

Lisa: "...I must be the only person who agreed [with the topic]."

Colleen (interrupting Lisa): "Wha - why did you agree?"

Lisa: "Because I'm a working woman. Because I've seen it. No, I'm not married but I work with -- all the women I work with are married and they all

have children and n-one in particular is having so many problems right now with trying to care for her children, and so y'know that's what I'm seeing everywhere with with working women is the nnnn is that society just right now — in fact I was discussing this with with one of the men I work with 'cause he's got the same problem with children — society right now is not set up in such a way that um both of the partners in a marriage can work, 'cause all you can do is send the children to school nnnnn babysitters; this guy's a former schoolteacher, so we were talking about the inadequac-inadequacies of the school as far as being a baby sitting system, y'know..."

Janet: "Yeah but that's not what they're really there for. It's like like OK there are problems with it, but still, like the men y'know it's like OK for a period of time — why can't the man be home? OK it's like why should it be — in my family my father's the one who's home; he's the one who cooks, does the laundry, and it's just ma-my mother goes out to work and so it's like why can't it be that the man is the one who's home for a few years? Let the wife get her career started off and have her do it while the children are young; the father's there part of the time and then maybe at some point then they switch back, the man is working all the time during the day and the woman either has — then can have her job because the kids are old enough to be in school and have it so that they can come home"

[Transcript of an experiment made at Cornell University, September 1978; some interruptions and noises of agreement have been omitted.]

With respect to **topic organization**, the connectedness and flow of sentence topics in Janet's paragraph illustrate the point. A good author would not, as Janet did, begin to state a point of view, suddenly interrupt herself to present an example, and then return to the original topic:

"...why can't the man be home OK it's like why should it be in my family my father's the one who's home; he's the one who cooks...so it's like why can't it be that the man is the one who's home for a few years..."

Similarly, Lisa's sentence topics appear in haphazard order:

- (a) WORKING WOMEN: she is a working woman and has seen it
- (b) COLLEAGUES: her female colleagues all have children
- (c) COLLEAGUE: a female colleague has a problem caring for her children
- (d) WORKING WOMEN: this problem is quite common with working women
- (e) SOCIETY: at present, society is not properly set up
- (f) COLLEAGUE: a male colleague has the same problem caring for his children
- (g) SOCIETY: in current society, school is the only place for some children

Under pressure, Lisa's **sentence organization** is poor, which explains her tortuous answer to Colleen: For example, from the outset, she is fully aware of the four ideas *children require care, currently, parents provide the only acceptable care, when parents work they can't provide care, and many mothers work outside home*. Her difficulty in answering arises from the inability to find phrases or words capable of expressing the multiple dependencies. So she tackles each part separately, giving specific examples; only after that is she able to state the relation between society, work, and child care:

...society right now is not set up in such a way that um both of the partners in a marriage can work, 'cause all you can do is send the children to school

And this only after she had tried and failed to say it once before:

...with working women is the nnnn is that society just right now —

Given more time, Lisa could have examined the relationships among the four concepts more carefully and used some standard ways of expressing the dependency *if not A then B* (such as "since not A, B" and "A, which causes B") in order to marshal the four topics as follows:

Since schools are not babysitters, the only place where children are properly cared for in today's society is at home, which causes a problem if both parents work

(It is much easier to produce polished text when you are writing at a terminal than when you are speaking into a microphone!) Alternatively, she might have found an interpretation of these topics that would provide her with a ready-made phrase. For example, this configuration of concepts is similar to, but not exactly the same as, the TOP (a high-level goal-plan configuration; see [Schank 82]) characterised by *if you want something done right, do it yourself*. This TOP would have been appropriate if the parents entrusted their children to a school and then at some later time were disappointed in the child care quality; however, Lisa's colleagues knew that schools did not provide good care from the outset. Had she interpreted the four ideas as an instance of this TOP, she may instead have said:

If you want your children brought up right, do it yourself; the schools are not going to look after your children properly. And if you're a working couple, there's nowhere else you can take your children in current society, so you have a problem

In addition, note in Lisa's and Janet's text the many instances of elucubration (a), false starts (b), rephrased sentences (c), and sentences in which crucial information is appended to the end (d):

- (a) *No, I'm not married but I work with — all the women I work with...*
- (b) *It's like like OK there are problems with it...*
- (b) *...but still like the men y'know it's like OK for a period of time —*
- (b) *it's just ma-my mother goes out to work*
- (c) *...that's what I'm seeing everywhere with with working women is the nnnn is that society just right now ---*
- (c) *and the woman either has — then can have her job...*
- (d) *we were talking about the inadequac inadequacies of the school as far as being a baby sitting system, y'know...*

When performing **topic selection**, Lisa and Janet both include examples. Lisa's two examples are essentially identical - why does she mention the male colleague after having

told of the female colleague with the same problem? Obviously she didn't check the second example well enough to ensure that it contained some new information. Note that this second example appears directly after her aborted attempt at making the general statement about society. It seems safe to assume that when she ran into trouble with that sentence, she needed some other topic to fall back on, and the extra reminding was the first thing at hand. Had she had more time, Lisa might have realized that it added nothing new to her point and decided not to say it.

The planning and realization tasks described here are not absolutely necessary for achieving the central goal to say (in some fashion or other) the main topic. In other words, the less pressure a speaker feels, the more time he has available to

- search memory for *additional sentence topics* under guidance of topic collection plans such as the CONVINCE plan described in chapter 4 (the opposite of this corresponds to blurting out whatever comes into your head).
- search for concepts (say, by running inferences, as described in chapter 3) that can be *interpretations* of the topic, and test the interpretations for appropriateness.
- search memory for instances similar to the topic to use as *examples* and select appropriate ones.
- determine the likely *affective effect* of the topic on the hearer in order not to make *faux pas*.
- try to juxtapose topics in affect-imputing *multi-predicate phrases*, as described in chapter 4.
- take time to select *appropriate phrases, sentence types, and words*.

When does a speaker feel time pressure? In an ideal world, the speaker will always have time to plan out his text completely, testing each topic, each phrase, each word for pragmatic suitability. This would satisfy the default speaker goals to present topics so that the hearer finds them intelligible and accepts their implications. However, another default

goal calls for the speaker to ensure that the hearer finds the conditions of the conversation acceptable; specifically, the speaker should not waste the hearer's time. Thus one factor that determines haste is the amount of time the hearer is willing to wait. In this regard, hearers have available a number of signals to speed up the tempo, such as saying "yes, yes" frequently, completing the speaker's sentences, or repairing his errors for him rather than allowing self-repair (see [Schegloff, Jefferson & Sacks 77, p 380]). On receiving such signals, the speaker should increase his desire to be quick. The opposite holds too: If the speaker speaks too quickly or spends too little time on each topic he violates the first default goal mentioned above; in such cases the hearer's signals are requests for clarification, and the speaker should slow down.

A number of other pragmatic factors influence the value of this rhetorical goal. The speaker must hurry when there is little time to complete the conversation, (for example, when he is making a long-distance telephone call). Relative social status plays a role too: when the speaker is subordinate to the hearer his level of haste must increase to help ensure that he will be entertaining. Based on these considerations, PAULINE's activation rules are (where the program's rhetorical goal **RG:haste** takes a value from the range *pressured*, *unplanned*, *somewhat planned*, *highly planned*):

1. set **RG:haste** to *pressured* if: the **time** is marked *little*, the **relative social status** is marked *subordinate*, and the **depth of acquaintance** is marked *acquaintances* or *strangers*
2. or set **RG:haste** to *highly planned* if: the **time** is marked *little*; or if **relative social status** is marked *subordinate* and the **depth of acquaintance** is marked *acquaintances* or *strangers*
3. or set **RG:haste** to *unplanned* if: the **time** is marked *much* and the **speaker-hearer depth of acquaintance** is marked *friends*
4. otherwise, set **RG:haste** to *somewhat planned*
5. then, reset **RG:haste** one step toward *pressured* if: hearer's **knowledge level** is marked *expert* and the **speaker's knowledge level** is marked *expert* or *student*

The degree of haste is one of the factors that makes spoken text different from written text. The prudent generator takes more time to plan when writing than when speaking, because a writer knows that (a) his text must be more grammatical than a speaker's, since it cannot contain incomplete sentences, retractions, and other mistakes, and that (b) it is much harder to recover from grammatical but uncompletable or foot-in-the-mouth sentences when they have been written than when they have been spoken or are still being planned.

5.4 Force

Forceful text is straightforward, direct, and has momentum. Generally, the effect is to draw the hearer's attention and to inspire him to action. Force is achieved by using a number of strategies that can be treated independently, for they each have characteristic effects that can be combined in other ways to produce different styles. (In the program, in fact, these component strategies are implemented separately as the rhetorical goals of style **RG:incitement**, **RG:aggression**, **RG:speaker-reference**, **RG:hearer-reference**, and **RG:warmth**.) However, for the sake of brevity, only the amalgam **RG:force** will be discussed here.

In order to determine the strategies of force used by speakers, a number of texts, taken from newspapers, advertisements, style books, speeches, and academic papers, were analyzed. The results can be summarized as follows: Aggressive text is calculated to capture and win the hearer's attention and to manipulate him by making him feel threatened or angry. It shares with inciting text the goal to exhort him to behave in a certain way, but where incitement is the carrot, aggressive text functions as the stick. (Much TV and radio advertising pretends to be inciting but is aggressive in this way.) In order to grip the hearer's attention, direct personal reference is common, both to the hearer and to the speaker. References to the setting are common too, as well as to circumstances and memories shared by the interlocutors. The following paragraph, (a), almost caricatures forceful style; compare it with the much quieter version (b):

(a) ...another disgrace. *The fact that America's homes are as filthy and messy as pigpens. Until now, I blamed those yappy women's libbers for the sorry conditions of our homes. They're the ones who encourage our gals to get out of the house and find jobs... But now I realize that a lot of the blame for crummy homes belongs to lazy housewives who watch soaps and The Price Is Right. These slobs could take a lesson from my wife Thelma Jean — the best little homemaker in America... I'm proud to say that on the day of the terrible space shuttle disaster, my little honeybun had no idea that anything special had happened. It wasn't until I called her that she knew. Thelma Jean wasn't glued to her TV set like so many of our lazy women.*

[*My America*, column by Ed Anger in **Weekly World News**, March 4, 1986]

(b) ...another disgrace, namely, the fact that America's homes are not clean and neat. *Although women's libbers may be responsible for the conditions of homes in this country, because they seem to be the ones who encourage housewives to get out of the house and find jobs, perhaps a lot of the blame for messy homes also belongs to less energetic housewives in whose homes TV is a constant distraction. Perhaps these women could learn from Thelma Jean Anger, an exemplary American homemaker, who never watches TV during the day.*

Based on the analysis, the strategies that a speaker or program must use to make its text forceful when it reaches decision points are:

- **(a) topic inclusion and organization:** enhance the sensitive aspects of topics by using appropriate phrases (such as the enhancer phrases described in chapter 4); include direct references to the interlocutors
- **(b) sentence organization:** make declarative and imperative sentences instead of requests or questions (for example, say "He doesn't know what he is talking about" rather than "It seems he doesn't know what he is talking about")

- **(c) sentence organization:** make short, simple sentences, by choosing not to link together sentences; by including at most one adverbial clause; by selecting short phrases where possible; by, during predicate construction, using the active voice
- **(d) sentence inclusion and word choice:** use enhancers — verbs, nouns, stress words — to stress affect-laden aspects of the topic
- **(e) word choice:** rather than flowery or unusual options, select forceful or simple, plain words and phrases

In contrast, in order to make text calm, the speaker must:

- **(a) topic inclusion and organization:** mitigate the sensitive aspects of topics by using appropriate phrases (such as the affect-imputing phrases) and euphemisms (such as "Uncle John has a bit of a headache after last night's party and will not come and play with you"); also avoid direct references to the interlocutors
- **(b) sentence organization:** make questions, requests, and signal opinions explicitly (using, for example, "I think", "it seems", "don't you agree?")
- **(c) sentence organization:** make longer, more involved sentences, by including many adverbial clauses; by conjoining and subordinating topics; by including words that can be elided; by, during predicate construction, using the passive voice
- **(d) sentence inclusion and word choice:** select phrases that denote a calm attitude (such as "well, well", "it would seem that", "let's see..."); mitigate sensitive aspects (using mitigators such as "only" and "just")
- **(e) word choice:** select long, flowery, less forceful words and phrases

As an illustration of the effect of these rules, note how these strategies are used in the following newspaper editorial:

Get the vote out to get Reagan out!

The time is at hand to end the presidency of Ronald Reagan. We must prevent four more years of policies dictated by the ultra-right.

If the people turn out to vote in record numbers we will have a new president and a new and better Congress. Every eligible voter ~~must~~ exercise his or her right to cast their ballot. That's the only way to assure a defeat for Reagan and Reaganism.

Everyone who has been affected by Reaganomics and Reaganism is needed to help mobilize the voting population of the United States to produce a record smashing, Reaganite smashing vote on November 6... The Reagan gang must cease to govern. They are union busters and wage cutters. They have cut Medicare and Social Security...

In the remaining hours of this campaign, and on Election Day, we urge our readers to

- *Saturate many shops with literature...*
- *Expand door-to-door cavassing...*
- *Arrange to take Election Day off from work...*

No sector of the electorate should be conceded to the Reaganites. Every vote must be fought for and won. It is critical that all victims of Reaganism not only vote, but also get out the vote. Start now!

[Editorial, Daily World (a pro-communist newspaper), November 1, 1984]

Clearly, this is forceful stuff. Sentences are short, simple, and declarative — the entire editorial contains only one conditional sentence and no questions. Stress words abound: “ultra”, “only”, “everyone”, “every”. Many words and phrases with punch appear: “dictated”, “exercise his or her right”, “urge”, “fought for”, “critical”, “victims”. After identifying which decisions give this editorial its forceful character, the corresponding inverted rules were applied, the resulting text is properly insipid:

Forceful text	Decision type	Insipid text
[The time is at hand] to end the presidency of Ronald Reagan. [We must] prevent four more years of policies [dictated] by the [ultra-right]. If the people turn out to vote in record numbers we will have a new president and a new and better Congress. Every eligible voter [must] [exercise his or her right] to cast their ballot. That's the [only] way to [assure] a defeat for Reagan and Reaganism.	d a c e e c d	[Now is a good time] to end the presidency of Ronald Reagan [in order to] prevent four more years of [right-wing policies]. If the people turn out to vote in record numbers we will have a new president and a new and better Congress. Every eligible voter [should] cast his or her ballot. [because] that's the way to defeat Reagan and Reaganism.
[Everyone] who has been affected by Reaganomics and Reaganism is [needed] to help mobilize the voting population of the United States to produce a record [smashing], [Reaganite smashing] vote on November 6. The [Reagan gang] must cease to govern. They are [union busters] and wage cutters. They have cut Medicare and Social Security.	e e b e d b e c e c b	[People] who have been affected by Reaganomics and Reaganism are [asked] to help mobilize the voting population of the United States to produce a record vote on November 6. [Isn't it time that] Reagan [and his staff] cease to govern. [since] they [oppose unions], cut wages, Medicare, and Social Security[?]
In the remaining hours of this campaign, and on Election Day, [we] [urge] our readers to...	e a b	In the remaining hours of this campaign, and on Election Day, our readers [are requested] to...
No sector of the electorate should be conceded to the Reaganites. Every vote must be [fought] for and won It is [critical] that all [victims of Reaganism] not only vote, but also get out the vote [Start now]	c e e e e d b c	No sector of the electorate should be conceded to the Reaganites. [therefore] every vote [should] be [worked] for and won It is [important] that all [who experience Reaganism] not only vote, but also [help to] get out the vote, [so please] start [doing so] now

When should a speaker make his text forceful? How does forceful text affect the hearer? Since it is short, direct, and has momentum, forceful text is appropriate when the speaker wants to inject energy into the hearer or the conversation -- thus, when he wants to incite him to action, draw him into the conversation, capture his attention, sway his opinions. In contrast, quiet text is placating and calm, therefore appropriate in more solemn or anguished occasions. PAULINE's rhetorical goal **RG:force** takes a value from the range *forceful, neutral, quiet*; it uses the following rules to find an appropriate value from its pragmatic setting:

1. set **RG:force** to *forceful* if: **desired effect on hearer's goals** is marked *activate*, that is, the speaker wants to activate a goal or plan in the hearer so as to affect his future behavior; or if **desire to involve hearer** is marked *involve*, that is, the speaker wants to spur the hearer to take part in the conversation; or if **desired effect on hearer's emotional state** is marked *anger*, for example if the speaker wants to goad the hearer; or if **atmosphere (tone)** is marked *informal*
2. set **RG:force** to *quiet* if: **desired effect on hearer's goals** is marked *deactivate*, that is, the speaker wants the hearer to forget about (deactivate) a goal or plan; or if **desired effect on hearer's emotional state** is marked *calm*; or if **atmosphere (tone)** is marked *formal*, for example, if the speaker wants to make the tone of conversation less boisterous; or if **desire to involve hearer** in the conversation is marked *repel*
3. otherwise, set **RG:force** to *neutral*

5.5 Summary

In summary, this section describes PAULINE's generation of a fictitious primary election between Carter and Kennedy in 1979 (from chapter 1) under five pragmatically different scenarios. The episode consists of a network of about 50 representation elements. After the five texts are listed, a short description follows of how the program acts in each phase of the generation process. Different program behavior results from different activated rhetorical goals; for purposes of illustration, the effects of only four rhetorical goals will be discussed: **RG:formality**, **RG:partiality**, **RG:detail**, and **RG:haste**.

In **case 1**, PAULINE must inform an acquaintance of the outcome of the primary and of the current status of both delegates. Neither interlocutor has opinions about the topic; both have the usual knowledge of the electoral process. This is achieved by giving PAULINE's characterizations of the speaker and the hearer their default values: normal interest in the topic, no sympathies or antipathies, calm emotional state, informal setting, normal conditions. In addition, PAULINE is given the following interpersonal goals:

- Hearer:
 - affect his knowledge -- *inform*
 - affect his opinions of topic - *no effect*
 - involve him in the conversation --- *no effect*
 - affect his emotional state - *no effect*
 - affect his goals *no effect*
- Speaker-Hearer Relationship:
 - affect hearer's emotion toward speaker *make like*
 - affect relative status *make equal*
 - affect interpersonal distance *make distant*

These values activate the following rhetorical goals: **RG:formality** is *colloquial*, **RG:partiality** is *impartial*, **RG:detail** is *details*, and **RG:haste** is *somewhat planned*. In this case, the program says

Case 1.

ON 20 FEBRUARY, CARTER AND KENNEDY WERE THE CANDIDATES IN A
PRIMARY IN MICHIGAN. CARTER LOST TO KENNEDY BY 1335 VOTES.
KENNEDY HAS A BETTER CHANCE OF GETTING THE NOMINATION THAN BEFORE
AT PRESENT. CARTER IS ALSO CLOSER TO GETTING THE NOMINATION THAN
BEFORE. BOTH CARTER AND KENNEDY WANT TO GET THE NOMINATION.

In **case 2**, PAULINE is sympathetic to Kennedy, while the hearer, the program's knowledgeable sibling, supports Carter. In this case the hearer is defined to have the knowledge state *expert*, with depth of acquaintance *intimate*, relative social status *equal*, and emotion *like*. The program has the goals to *make close* the interpersonal distance and to *inform* the sibling. Both interlocutors have the same sympathy: Carter. These values activate the rhetorical goals as follows: **RG:formality** is *colloquial*, **RG:partiality** is *implicit*, **RG:detail** is *details*, and **RG:haste** is *somewhat planned*. In this case, PAULINE says

Case 2.

WELL, SO CARTER LOST THE PRIMARY TO KENNEDY BY 1335 VOTES.

Case 3 is similar to case 2, but the hearer is a *friend* and social *equal* (say, a colleague) who is not as expert as the sibling (i.e., knowledge level is *student*). But now both interlocutors have opinions: PAULINE's sympathy is for Kennedy and the hearer's for Carter. The conversational tone is still *informal*, which gives PAULINE the time to prepare its text well, in order to try to convince Bill that Kennedy is going to win. The rhetorical goal **RG:formality** takes the value *colloquial*, **RG:partiality** is *implicit*, **RG:detail** is *all* (details and interpretations), and **RG:haste** is *planned*. The program says

Case 3.

KENNEDY DIMINISHED CARTER'S LEAD BY GETTING ALL OF 21850 VOTES IN

THE PRIMARY IN MICHIGAN. IN A SIMILAR CASE, CARTER DECREASED UDALL'S LEAD IN A PRIMARY IN 1976, AND HE EASILY TROUNCED UDALL TO BE NOMINATED BY 2600 DELEGATES. I AM REAL GLAD THAT KENNEDY IS NOW CLOSER TO GETTING THE NOMINATION THAN BEFORE.

In **case 4**, PAULINE is a Carter supporter and is speaking formally — say, making a speech at a debate — so that the conversation time is only *some* and the tone is *formal*. The audience is presumed to support Kennedy, while the program supports Carter. Thus PAULINE has the goal to *switch* the hearers' opinions of the topic, to make them *respect* the speaker, and yet feel *distant*. These values activate the rhetorical goals as follows: **RG:formality** is *highfalutin*, **RG:partiality** is *explicit*, **RG:detail** is *details*, and **RG:haste** is *somewhat planned*. In this case, PAULINE says

Case 4.

I AM PLEASED TO INFORM YOU THAT CARTER HAS IMPROVED HIS CHANCES OF WINNING THE NOMINATION. AT THE PRESENT TIME, CARTER HAS MANY MORE DELEGATES THAN HE HAD IN THE PAST; ALSO, CARTER HAS MANY MORE THAN KENNEDY DOES.

Finally, in **case 5**, PAULINE is a Carter supporter and is speaking to its boss, an irascible Kennedy man. They are making a long-distance telephone call, which gives the program *little* time and makes conversational conditions *noisy*. Furthermore, the program is *distant* from its boss, does not wish to anger him (desired emotional effect is *calm down*), and still wants to make him feel socially *dominant*. The four rhetorical goals get the following values: **RG:formality** is *colloquial*, **RG:partiality** is *implicit*, **RG:detail** is *interpretations*, and **RG:haste** is *pressured*. To its boss, the program says

Case 5. ...

...nothing!

Note that PAULINE does not generate widely different versions of the central topic. To do so would be easy: the program would simply have to discriminate to one of a number

of greatly different sentence forms (however they are represented) and then fill it in. But this would prove nothing beyond the fact that PAULINE uses relevant pragmatic aspects in its discrimination process. In this work, the question is more subtle: *how is additional information implicitly encoded in text?* — in other words, how can the same phrases and words be selected, rearranged, and juxtaposed in order to convey different information? When this question has been answered, the correct way of treating a number of greatly different sentence forms will be easy to solve too.

5.5.1 Topic Collection

In all five cases, PAULINE is given a single input element that represents the primary election between Carter and Kennedy as the central topic of conversation. (This is done for purposes of comparison; any of the 50 representation elements required for the story could have been given as input topic(s), but the resulting texts would, of course, be different.) From the central topic, PAULINE can search for relevant additional sentence topics. As described in chapters 2 and 4, the program has three topic collection plans: RELATE, DESCRIBE, and CONVINCe.

Using the selection rules described, the RELATE plan is used in case 1 (speaking to an acquaintance), since the interlocutors do not hold different affects for the topic, and the CONVINCe plan is selected in the other four cases. The plan's steps are applied to the representation in order to suggest candidate topics. In case 5, speaking to the boss under *pressured RG:haste*, PAULINE has to say each candidate topic as soon as it is found; in the other cases, it can apply all the plan steps and collect a number of candidate topics before proceeding with further planning. In cases 2 and 3, the candidates are the topics that support a pro-Kennedy argument, namely Kennedy's and Carter's outcomes and Kennedy's current delegate count. In cases 4 and 5, the collected topics are Carter's delegate count, including the facts that it is larger than it was before and still is larger than Kennedy's.

5.5.2 Topic Organization

After it has collected candidate topics and before it says them, PAULINE performs a number of topic organization tasks (described in more detail in chapter 7): it orders the topics (as described in chapter 7); checks whether the topics can be appropriately interpreted (as described in chapter 3); it checks whether sentences of explicit opinion (such as "I am angry about that") should be included (as described in chapter 2); it tries to juxtapose topics in multi-predicate phrases (as described in chapter 4).

When PAULINE speaks to its boss (case 5), the candidates it collects all oppose his sympathies. Since **RG:partiality** is *implicit*, the program is required to mitigate such sensitive topics (say, by using appropriate multi-predicate phrases or using an interpretation that subsumes them). However, these are time-consuming tasks, and the program's rhetorical goal **RG:haste**, with value *pressured*, does not permit PAULINE to do more than test the candidates for affective suitability; hence, in this case, PAULINE cannot say *anything!* (The first time PAULINE did this, I thought it was a bug in the program.)

In case 3, however, PAULINE has more time to perform the planning tasks. In particular, the collected facts (Carter was and still is ahead though Kennedy won the primary) match the pattern defining the interpretation *narrow lead*. Indexed under the interpretation, as described in chapter 3, the program finds two reminders: Hart narrowing Mondale's lead in 1984 (but still losing the nomination), and Carter narrowing Udall's lead in 1972 (and eventually winning). Since **RG:detail** is not set to *details*, the program is allowed say the interpretation; also, the value *planned* for **RG:haste** allows time to select the appropriate reminding (by mapping the equivalent role fillers and checking affects) and to cast all this in suitable phrases. (In addition, the newly created interpretation is added to memory: when PAULINE tells the example again the interpretation is immediately found and can be said directly.)

In case 2, PAULINE's goal **RG:detail** calls for low-level details, and thus it doesn't search for interpretations but simply says the two outcomes. Similarly, in case 1, the program organizes its details *impartially* by alternating topics with opposing affects. And

in case 4, since it doesn't find interpretations off the collected Carter-supporting topics, the program simply orders them and casts them into multi-predicate phrases.

5.5.3 Sentence Organization

When realizing a sentence, PAULINE must select the subject, select which adverbial clauses to say before the subject, select a verb that doesn't require pragmatically sensitive aspects (such as "win", with no direct object rather than "beat"), and order the predicate clauses. Furthermore, it must select appropriate aspects to say as adverbs and adjectives, build noun groups, and select appropriate words.

In cases 2, 3, and 4, the strategies for partial text (as described in chapter 4) cause PAULINE to include the clauses "I am glad that", "I am pleased to inform you", and the affective adjectives and stress words "many", "all of", and "easily". In case 4, when PAULINE formally addresses the hostile audience, the strategies for formal text outlined above cause it to use formal wording in "I am pleased to inform you" and "at the present time", to place the latter before the sentence subject instead of after it, as well as to include the extra verb "does". To produce explicit and implicit partial text, the program selects nouns and verbs that carry affect; formal words and phrases help achieve the goal of **RG:formality** in case 4.

5.5.4 Analysis

In summary, the following tables illustrate the effects of the rhetorical goals on the examples:

Case 1 (to an acquaintance): *colloquial, impartial, details, somewhat planned*

text	decision	rhet. goal value
Topic: central topic	RELATE plan	
[] CARTER AND KENNEDY WERE THE CANDIDATES IN A PRIMARY [IN MICHIGAN] [ON 20 FEBRUARY].	no clauses before clauses after subject	<i>colloquial</i> <i>colloquial, planned</i>
Topic: result	RELATE plan	
CARTER [LOST] TO KENNEDY BY [1336] VOTES.	neutral verb neutral details	<i>impartial</i> <i>impartial, details</i>
Topic: outcome with good affect for Kennedy	RELATE plan	<i>impartial</i>
AT PRESENT, KENNEDY HAS A BETTER CHANCE OF [GETTING] THE NOMINATION THAN [] BEFORE.	 informal word elide <i>he had</i>	 <i>colloquial</i> <i>colloquial</i>
Topic: outcome with good affect for Carter	RELATE plan	<i>impartial</i>
CARTER IS ALSO CLOSER TO [GETTING] THE NOMINATION THAN [] BEFORE.	separate sentence informal word elide <i>he was</i>	<i>colloquial</i> <i>colloquial</i> <i>colloquial</i>
Topic: actors' goals (twice)	RELATE plan	
BOTH CARTER AND KENNEDY [WANT] TO [GET] THE NOMINATION.	informal verb informal verb	<i>colloquial</i> <i>colloquial</i>

Case 2 (to an expert sibling): *colloquial, implicit, details, somewhat planned*

Topic: results with good affect for Kennedy	CONVINCE plan	<i>implicit, planned</i>
[WELL, SO] CARTER LOST THE PRIMARY TO KENNEDY BY [1336] VOTES.	informal, to sibling details	<i>colloquial</i> <i>details</i>
Topics: outcomes with good affect for Kennedy	CONVINCE plan	<i>details</i>
	suppressed due to hearer knowledge	

Case 3 (to a friend): *colloquial, implicit, all (details and interpretations), planned*

Topic: results with good affect for Kennedy	CONVINCE plan	<i>implicit</i>
[] KENNEDY	no clauses before	<i>colloquial</i>
[DIMINISHED] CARTER'S [LEAD]	interpretation	<i>all, planned</i>
BY [GETTING]	informal verb	<i>colloquial</i>
[ALL OF]	stress word	<i>implicit</i>
[21850] VOTES	details	<i>all</i>
[IN THE PRIMARY] [IN MICHIGAN].	clauses after subject	<i>colloquial</i>
Topic: reminding	indexed off interp	<i>planned</i>
IN A SIMILAR CASE, CARTER DECREASED	reminding	<i>implicit, planned</i>
UDALL'S LEAD IN A PRIMARY		
IN 1976, AND HE [EASILY]	stress word	<i>implicit</i>
[TROUNCED] UDALL TO BE NOMINATED	stress verb	<i>implicit</i>
BY [2600] DELEGATES.	details	<i>all</i>
Topic: outcome with good affect for Kennedy	CONVINCE plan	<i>implicit</i>
[I AM REAL GLAD THAT]	informal opinion	<i>colloquial, explicit</i>
KENNEDY IS [NOW] CLOSER TO	clause after, informal	<i>colloquial</i>
[GETTING] THE NOMINATION THAN	informal verb	<i>colloquial</i>
[] BEFORE.	elide <i>he was</i>	<i>colloquial</i>

Case 4 (making a speech): *highfalutin, explicit, details, somewhat planned*

Topic: results with good affect for Carter	CONVINCE plan	<i>explicit</i>
---	none	
Topic: outcome with good affect for Carter	CONVINCE plan	<i>explicit</i>
[I AM PLEASED TO INFORM YOU] THAT	formal opinion	<i>highfalutin, explicit</i>
CARTER HAS [IMPROVED HIS CHANCES]	formal phrase	<i>highfalutin</i>
OF WINNING THE NOMINATION.		
Topic: outcome with good affect for Carter	CONVINCE plan	<i>explicit</i>

[AT THE PRESENT TIME], CARTER HAS	clause before, formal	<i>highfalutin</i>
[MANY] MORE DELEGATES THAN	stress	<i>explicit</i>
[HE HAD] IN THE PAST;	no elision	<i>highfalutin</i>
[ALSO], CARTER HAS	long sentence	<i>highfalutin, planned</i>
[MANY] MORE THAN	stress	<i>explicit</i>
KENNEDY [DOES].	no elision	<i>highfalutin</i>

Case 5 (to the boss): *colloquial, implicit, interpretations, pressured*

Topic: results and outcomes for Carter	CONVINCE plan	<i>implicit</i>
...	no time for mitigation	<i>pressured</i>

5.6 Conclusion

In chapter 1, we asked the question “why and how is it that we say the same thing in different ways to different people, or even to the same person in different circumstances?”, and suggested that the answer to the “why” part of the question relates to the communication of non-literal information, especially pragmatic information, in the text. This chapter described some of the answer to the “how” part. In doing so, it followed an algorithmic approach to the creation of style in language. Though many of the program’s rules for achieving rhetorical goals of style may need refining, the underlying claim — that style is the result of following a coherent policy when making decisions during the process of generation — is undeniable.

Chapter 6

Grammar and a Phrasal Lexicon

Abstract

The question addressed in this chapter is: how should language be represented in a generator? In particular, how do the concepts the generator must express, the grammar it is to use, and the words and phrases with which it must express them, relate? The argument put forth here is that all the structural aspects of language — rules of grammar, phrases, word patterns — should be represented together in the lexicon in the form of phrases and features of words, and that the lexicon should be closely bound to the system's conceptual network.

6.1 The Three Tasks of a Generator

As input, language generation programs are given sets of representation elements. (Assume that the generator's input consists of structures built using a representation scheme that is not based on the syntax of any language. One example of such a scheme is Conceptual Dependency Theory, [Schank 72, 75], extended in [Schank & Abelson 77] and [Schank 82]; a similar scheme is developed in [Jackendoff 85]. Assume also that the representation structures are defined within a property inheritance network such as those in common use,

as described, for example, in [Stefik & Bobrow 85], [Charniak, Riesbeck, & McDermott 80], or [Bobrow & Winograd 77].) If an element can be said directly, in one word or frozen phrase, the generator's task is easy; otherwise, the generator has to break up the element into parts and concentrate on each part, recursively. The order in which it examines the parts will be reflected in the order of the words of the text; hence, the generator must use the ordering conventions of the language to guide its traversal of the input.

During its traversal, the generator must consider progressively "smaller" pieces of the input element, or it must consider pieces from a progressively "narrower" point of view, so that it will eventually produce text and not just blindly continue traversing the whole network in which the elements are defined. That is to say, if the generator starts out with the goal of making a sentence about some input *X*, then its next goal could be to make a sentence subject of some part *X1* of *X*, and its next goal to make a noun phrase of some part *X2* of *X1*, and so forth. This sequence of goals must eventually end in the "narrowest" goal, namely the goal to output one or more words from the lexicon without spawning any further goals. Though this sequence need not monotonically decrease in scope (since, for example, whole sentences can be relativized and subordinated to other sentences), it must always terminate. (Of course, the pieces of the input do not really become "smaller", however size is measured; it is simply convenient to think of them doing so, in the sense that the agent of an event is somehow contained within the event and that the agent's age is in turn contained within the agent, which means that more words of the sentence describe the event than the agent, and in turn more words the agent than the age.)

Therefore, from the piece of the input under consideration, the generator must select the following: which section(s) it is going to work on next; in what order it is going to do so; and what work it is going to do on each section; — in such a fashion that it is guaranteed eventual termination. Thus the generator must perform three tasks:

- **inclusion:** select which portions of the input to consider further (which portions will eventually appear in the text)
- **ordering:** select the order in which to consider them (in which order they will appear in the text)

- **casting:** select a syntactic class or environment in which to say each portion (in which form they will eventually appear in the text)

Consider the following example, using a representation similar to Conceptual Dependency notation ([Schank 72, 75]), where MTRANS means "transfer of information":

```
#(ACTION = MTRANS-6
  [ACT : MTRANS]
  [ACTOR : JIM]
  [OBJECT : #{STATE-CHANGE = DEATH-10
    [TYPE : HEALTH]
    [ACTOR : JANET]
    [FROM : ALIVE]
    [TO : DEAD]]}]
  [FROM : JIM]
  [TO : SUE]
  [MANNER : QUIET])}
```

From this representation, a generator should be able to produce at least the following sentences:

1. Jim told Sue that Janet died
2. Jim told Sue of Janet's death
3. He told her of Janet's dying
4. Jim told Sue
5. He whispered to her that Janet died
6. Jim quietly told Sue of it

7. Quietly, Jim let Sue know that Janet died
8. Jim whispered
9. Sue was told by Jim that Janet died
10. Sue heard of Janet's death
11. She heard of Janet's dying from him
12. Sue was quietly informed of Janet's death
13. Janet's death was what Jim quietly informed Sue of
14. That was what Sue heard from Jim

These sentences are produced in the following way: Initially, of course, the generator simply has the goal to make a sentence from MTRANS-6. Its first decision is: which aspects should be included? JIM and SUE and DEATH-10? Only JIM and SUE? Only DEATH-10? Then, if more than one are selected, it has to choose a sentence subject (in sentences (1) through (8), JIM is the subject, and in (9) through (12), SUE is). It also has to decide whether to include the adverb QUIET and how to order it with respect to the rest of the sentence (compare (5), (6), and (7)). When it starts building the subject, the generator must make a casting decision -- actions and state-changes must be cast as nominals or pronominalized (DEATH-10 in (13) and (14)); objects can be named, described or pronominalized ((1) and (3)). Later, when it builds the predicate, inclusion decisions pertain to adverbs ((5) and (6)) and to other parts of the topic ((10) and (11)); casting decisions include verb choice ((1), (5), (10)) and predicate form (see (1), (2), and (3)).

The form of each generated sentence is determined by the sequence of inclusion, ordering, and casting decisions made in the realization process. At any point in the process, the generator needs information on which linguistic options exist -- which decision tasks it must/may perform on the current input. The question of interest in this chapter is *how and where in the generator should this information reside?*

6.2 Formative Information

Most work on the representation of the structure of language makes a distinction between the *grammar* and the *lexicon*. The former is a body of the rules that govern how words can be put together; the latter is the collection of words and their idiosyncratic features. See, for example, [Chomsky 65, p 84] (his italics):

The grammar will contain no rules...that introduce the formatives belonging to lexical categories. Instead, the base of the grammar will contain a *lexicon*, which is simply an unordered list of all lexical formatives.

In this spirit, most generator programs contain as distinct entities: a set of grammar rules, a lexicon, and a mechanism that produces text (by accepting an input representation, building a syntactic tree structure on applying the rules of grammar to the input, inserting into the tree lexical entries that are accessed from the input representation, and finally saying the words).

6.2.1 Rules of Grammar

The rules of English grammar are concerned with what can be called syntactic environments. Based on their function, these rules can be divided into two groups. Rules from one group specify the *order of environments* within encompassing environments – for example, within the environment NOUN GROUP, the order

[ARTICLE ADJECTIVES HEAD-NOUN POST-NOMINAL MODIFIERS]

or, within a PREDICATE environment, the order of various noun groups:

[NG (subject) VERB NG (object) NG (location, direction, etc.)]

Rules from the other group specify *how different environments and their relationships are signalled* – for example, the case information provided by “’s” in “John’s book” or by the preposition in “to the store”; or the number agreement between subject and verb.

Some attempts at writing rules of the first kind — the ordering or *formative* rules — don't take any actual words into account at all. Words are simply inserted during the generation process (see, for example, [Simmons & Slocum 72]). But divorcing the formative rules from the lexical entities can cause problems. For example, generators built along these lines run the risk of building a syntax tree into which they cannot grammatically insert words, as in (b):

- (a) John beat Pete in the race
- (b) * Pete lost John in the race (*meaning*: Pete lost the race to John)

To ensure that this doesn't happen, you must either make the rules of grammar smart enough to distinguish between such cases as subjects that win and subjects that lose, or you must associate the various sentence structures with the words that control them (such as "beat" and "lose"), and make the rules examine the words in order to build appropriate trees. Obviously, the first alternative amounts to building rules that depend on words in any case, so in most systems much information about how words can combine with other words is associated with the words themselves. Thus some words in the lexicon carry features that constrain their combination with other words and syntactic environments. However, these features do not provide complete formative (ordering) information. For example, in the lexicon in [Stockwell, Schachter & Partee 72], the word "let" includes the features

"let":
 + V
 ADJ
 + TO-DEL
 + [--- +NEUT +DAT LOC INS +AGT]
 + DAT → OBJ

(that is, "let" is a verb but *not* an adjective (which some linguists consider a type of verb); "to" is deleted (otherwise, "John let Pete to win the race"); the predicate may not contain

the cases INSTRUMENTIVE or LOCATIVE .) Still, these constraints are not sufficient to prohibit sentences such as "John let win the race Pete" or "win the race let John Pete".

Associating this grammatical information with individual words subverts part of the function of formative grammar rules into the lexicon. The subversion takes place to varying degrees. In the transformationalist approach (see, say, [Chomsky 57, 65] or [Stockwell, Schachter & Partee 72, p 719]), the generator accesses the lexicon twice: once (after applying the phrase structure rules that build the basic sentence pattern, but before the transformation rules that reorganize it) for the words, such as verbs, with information used by the formative grammar rules, and once again (at the end, just before realization into speech or writing) for the words without this information, such as prepositions and pronouns.

In other approaches, the subversion is stronger: the definitions of words also include ordering information explicitly. This is the case in the systemic/functional tradition, as embodied by systemic grammar [Halliday 76] and [Mann 82], functional grammar [Kay 79], unification grammar [Kay 84], lexical functional grammar [Kaplan & Bresnan 83], and the grammar developed by [Gross 84] and [Danlos 85]. For example, the verb "beat" contains the formative pattern:

"beat":

[VERB [OBJECT loser (mandatory)] [PREPGROUP instance *in*]
[DIFFERENCE *by*]]
"She [beat] [him] [in the election] [by 3 votes]"

Here *loser* and *instance* indicate which aspects of the input element to say in the environments. In comparison, the verb "win" contains:

"win":

[[VERB [OBJECT instance] [DIFFERENCE *by*]]
"She [won] [the election] [by 3 votes]"

and "lose" contains:

“lose”:

[[VERB [OBJECT instance] [PREPGROUP winner *to*] [DIFFERENCE *by*]]

~He [lost] [the election] [to her] [by 3 votes]~

This method works well for cases where words — typically, verbs and nouns — require idiosyncratic combinations of words. Whenever the generator encounters a word with formative information, it uses that information to help build its sentence. But what of the general formative rules that are not tied to specific words? For example, in unification grammar, the functional description for NOUN PHRASE is

CAT = NP

PATTERN = (... N ...)

either: [ADJ = NONE]

or: [PATTERN = (ADJ ...)]

[ADJ = [CAT = ADJ]]

[[LEX = ANY]]

either: [PP = NONE]

or: [PATTERN = (... PP)]

[PP = [PATTERN = (PREP NP)]]

[[CAT = PP]]

[[PREP = [CAT = PREP]]]

[[[LEX = ANY]]]

[[NP = [CAT = NP]]]

(Reading from the top, this means: a NOUN PHRASE must have a NOUN. It need not have adjectives, but if it does, they are ANY words of CATEGORY ADJ and precede the noun. The noun phrase need not have any preposition phrases either, but if it does, they are of CATEGORY PP and follow the noun. Here, PPS consist of a PREP, where the preposition is ANY word of CATEGORY PREP, and an NP. This grammar is used in the generators of McKeown

[McKeown 82] and Jacobs [Jacobs 85], and the latter mentions some implementational difficulties and proposes solutions to them.) Rules such as these have but one function: to provide the types and the order of the constituents of syntactic environments. But this is exactly the function of the formative patterns associated with verbs and nouns, as described above! From a functional perspective, there is no reason why general formative rules of grammar should be viewed as being different from the formative patterns contained in the lexicon. They serve the same purpose. Therefore, they should be defined and used in the same way as the verb patterns are. Then, though not associated with specific verbs and nouns, all the standard phrase structure entities (S, VP, etc.) can be incorporated into the lexicon and accessed in a unified manner.

6.2.2 Lexical Entries

With respect to the lexicon, linguists have held the following positions:

The transformationalist hypothesis [Chomsky 65]: lexical entries contain *features*, functional primitives that are interpreted by grammatical rules to determine how they fit into the syntactic environments of sentences. (Sentences are built up by rules that transform the basic sentence pattern, step by step, into questions, complex sentences, etc.) As described above, the lexicon is accessed twice. The features of words already included in the tree may be modified by some transformation rules to create derived words, such as "proposal" from "propose".

The lexicalist hypothesis [Stockwell, Schachter & Partee 72]: similar to the above, but the lexicon contains parallel but distinct entries for words such as "propose", "proposal".

The deep case hypothesis [Fillmore 68]: the syntactic environments of sentences are determined by functional primitives called *cases*. These cases are attached to certain lexical entries, which also contain features that determine their behavior in syntactic environments. Transformation rules do not exist.

In [Fillmore 71], Fillmore states that the lexicon of a generative grammar must make available to its users, for each lexical item, at least

1. the nature of the deep-structure syntactic environments into which the item may be inserted;
2. the properties of the item to which the rules of grammar are sensitive;
3. the presuppositions or 'happiness conditions' for the use of the item, the conditions which must be satisfied in order for the item to be used 'aptly';
4. its meaning; and
5. the phonological or orthographic shapes which the item assumes under given grammatical conditions.

The argument for incorporating the formative rules of grammar into the lexicon, as discussed above, would add to these requirements the following

6. the order of the syntactic environments required by the item

6.2.3 Becker's Phrasal Lexicon

There is a lot more to language than grammar and words. Though not discussed as much by linguists, frozen and partially frozen phrases must also appear in the lexicon; generators must be able to use them to create sentences in the same way it uses grammar rules and words. This view is engagingly described by Becker in [Becker 75] (from the abstract):

...[U]tterances are composed by the recitation, modification, concatenation, and interdigitation of previously-known phrases consisting of more than one word. I suspect that we speak mostly by stitching together swatches of text that we have heard before... A high proportion of utterances are produced in stereotyped social situations, where the phatic and ritualistic functions of language demand not novelty, but rather an appropriate combination of formulas, clichés, allusions, slogans, and so forth.

Becker estimates that we know about as many stock phrases as we know single words (about 25,000), and about as many lexical similes (such as "pleased as punch" or "white as a sheet") as there are strong verbs in English (some 100). Thus any study of language that limits itself only to words and ignores phrases is hopelessly incomplete. His categorization of the types of phrases is (p 6):

1. **Polywords:** multi-word fixed phrases that are interchangeable with single words, such as: *for good; two bits; the facts of life*
2. **Phrasal Constraints:** short patterns (part of a sentence) that express an idea and allow some variability, such as: *by [pure, sheer] coincidence; [kick, kicked] the bucket*
3. **Meta-Messages:** phrases that link to previous conversation to indicate the train of thought, such as: *for that matter; guess what?; it's simply not the case that*
4. **Sentence Builders:** patterns for sentences expressing a stereotypical scenes, such as: *X gave Y a song and dance about Z; sell X short; X on a wild goose chase*
5. **Situational Utterances:** complete utterances matching specific circumstances, such as: *how can I ever repay you?; no way!; that's a good question*
6. **Verbatim Texts:** memorized texts, such as: the Pledge of Allegiance; *not by the hair of my chinny-chin-chin!*; advertisements

The verb- and noun-based formative patterns discussed above exist in the lexicon in addition to these phrases. As Becker says (p 32):

...most of the lexical phrases that we actually use are too humble and uninteresting that they would never appear on a list devoted to picturesque expressions like *Davey Jones's Locker*. Yet these humble expressions do most of the work of language production for us.

In this view of language, then, the rules of grammar that deal with the ordering of words and syntactic environments are contained in the lexicon as frozen, semi-frozen, and very general phrases. The lexicon is the sole repository of the patterns that make up language – some very specific, some very general. With respect to formative function, no principled difference exists between general patterns such as [SUBJECT VERB OBJECT] and specific ones such as [*the facts of life*], since intermediate patterns exist along the whole range of generality — patterns such as [[VERB *bury*] *the hatchet*] (giving “buried the hatchet” and “will bury the hatchet”, but not “bury the hatchets” or “bury the red hatchet”). Thus multi-predicate phrases (such as the enhancer phrase

[*not only* [SENTENCE (verb relocated, with “do”)] . *but* [SENTENCE]]

described in chapter 4) exists in the lexicon side by side with the verbs “beat”, “win”, and “lose”, discussed earlier. (The use of such phrases in language analysis is described in [Wilensky 81] and [Riesbeck & Martin 85], and some research on how they may be learned is reported in [Zernik & Dyer 85].)

Just as for verbs and nouns, phrases in the lexicon are linked to the concepts in memory that they describe. If an idiosyncratic phrase exists for the expression of a memory concept, the generator must have the option of using it instead of general all-purpose sentence-formation rules. Though there is no reason to associate the rule [SUBJECT VERB OBJECT], or even the verbs “huff” and “puff”, specifically with the story of the three pigs, the fixed phrase “he huffed and he puffed” belongs just there. Similarly, “kick the bucket” is tied to the state change DIE; “the big apple” is tied to New York City. And *nothing prohibits specific phrases from being ungrammatical*. This fact makes it impossible in principle to capture all the forms of language in a few general rules; thus, the linguistic endeavor of trying to create a formal, complete, consistent set of rules to describe all of language tries to describe structure where there is none, and is therefore destined to fail.

Thus, in summary, verbs, nouns, and other words in the lexicon are associated with the representation elements they describe. When idiosyncratic forms of expression exist, the formative patterns are associated with the lexical entries. Similarly, fixed and semi-fixed phrases describe representation elements; they are associated with representation elements in the same way; and they have similar formative patterns. Formative patterns are built

out of other patterns and lexicon entries. Although the general formative rules of grammar are not associated with any specific representational element, their formative patterns are defined similarly. This homogeneity enables the generator builder to add new forms of expression — words, phrases, or rules of grammar — with ease.

6.3 Syntax Specialists

6.3.1 PAULINE

Many generators rely on some central process to examine the input representation, to check its features, and to perform the inclusion, ordering, and casting of its aspects. This approach is most practical when the three decision types are relatively straightforward. This is the case in most functional/systemic generators to date: typically, the inclusion decisions simply are of the form “does aspect X appear in the functional description (the formative pattern)?”; ordering is given by the pattern; and casting is given by the pattern and by the feature constraints of the parts of the input. However, generators that are able to realize the same input in various ways (say, by taking into account pragmatic issues, as PAULINE does) have to make more complex decisions. In addition to syntactic constraints, their inclusion decisions depend on the pragmatic import of the pieces of the input; their ordering decisions, where alternatives exist, may carry pragmatic weight; and pragmatic issues can affect how pieces are cast as well. And, just as it makes sense to associate information about idiosyncratic syntactic phenomena with the words that control them, it makes sense to associate the pragmatic decisions with words (and other lexicon elements) as well. *It is sensible to encode all the relevant decisions in the functional descriptions of the lexicon entries themselves; that is, to spread the functionality of the central process into the lexicon.*

In PAULINE, the extended functional descriptions are called *syntax specialists* — each syntactic goal is achieved by a procedure, the specialist, that accepts a piece of input, performs the three tasks, and produces an ordered list of words and/or other syntactic goals,

each associating another specialist with a piece of the input. Thus syntax specialists are the repositories of information about the linguistic options. They control the performance of the three types of decision. Sometimes the specialists are very simple — so simple that they contain no procedural information — and then they are implemented as patterns. Alternatively they may be quite complex — directing much processing and altering the state of the generator — and then they are implemented as procedures.

The specialists correspond to the *clause templates* of [Danlos 85, 87], or to the *realization classes* of McDonald's generator MUMBLE (for example, [McDonald & Pustejovsky 85]); they can be viewed as implementations of the *systems* in the systemic grammar of [Halliday 76] (for a clear exposition, see [Patten & Ritchie 87]), and so resemble the systems in Nigel, the systemic grammar implemented by [Mann 82, 83a, 83b]. However, PAULINE's specialist functions differ from systems in a number of ways — most importantly, they are not activated whenever their input conditions are fulfilled, but rather are activated in a sequence determined by their predecessors (described at the end of the paper). This is a simplification of the systemic scheme. PAULINE's specialists differ from Nigel's systems in particular in that they can index phrases as well as words in the lexicon: they can index to more than one word via discrimination nets, as is described below; and their decisions refer to pragmatic criteria as well as to grammatical criteria (although Nigel makes provision for (unimplemented) pragmatic criterion functions called *choosers* [Mann 83a]).

Each syntax specialist must achieve the goal to create its syntactic environment with the input it receives. Thus the generator's NOUN GROUP and RELATIVE CLAUSE specialists make different decisions when given the same input element. Starting with the representation of "John shot Mary with a gun", when the generator's goal is to make a SENTENCE, it can say that sentence; if its goal is to make a noun group, the NOUN GROUP specialist can return "John's shooting of Mary with a gun"; and the RELATIVE CLAUSE specialist may produce "that John used to shoot Mary" or "who shot Mary with a gun". Each specialist must know of the different ways its goal can be achieved, and must be able to select an appropriate alternative. For example, an alternative noun group formulation for the above example is "John's use of a gun to shoot Mary"; and an alternative relative clause is "who was shot with a gun by John". The criteria by which these decisions are made can

be grouped into three classes: syntactic, semantic, and pragmatic. *Syntactic criteria* are binding: if they are ignored, ungrammatical sentences result. (For example, when saying a verb, choices concern singular or plural endings, appropriate tenses, and aspect.) *Semantic criteria* depend on the nature of the input and its relations to other concepts and the constraints of use of words. (For example, for the representation elements INGEST, the verb must match features of the OBJECT: a liquid gives "drink", a solid "eat", and a gas "breathe". This idea was first described in [Goldman 75].) Clearly, some input representations may be handled by a number of such syntax specialists. Picking one can be a problem. In this regard, pragmatic criteria can help make the decision¹. *Pragmatic criteria* relate to the affective values of words and their interactions with the speaker's goals, and are determined by strategies such as those described in chapters 3 to 5. The responsibility for accessing the relevant syntactic, semantic, and pragmatic information resides within each specialist.

6.3.2 Specialists and Phrase Structure Symbols

It is quite natural to identify certain syntax specialists with their equivalent phrase structure entities: to claim that when a linguist says "there exists a grammatical concept called noun group" he means that speakers have a collection of neurons that act in the brain in the way the noun group syntax specialist acts in a generator.

This identification can help answer some linguistic questions. For example, at issue for a number of years has been the question whether certain SVO languages are configurational

¹In [McCawley 78], McCawley notes that additional information is conversationally implicated, under the Gricean cooperative principle, when the speaker chooses a less direct way of saying something than the most straightforward one, if such exists. For instance, he uses Householder's example ([Householder 75], p. 75) that "pale red" is not "pink", whereas "pale blue", "pale yellow", "pale green", etc., all correspond to {color + white} in the color wheel. Similarly, in contrast to "let me in", "let me come in" implies that the speaker doesn't want to partake in the activities inside. This point indicates that the generator's casting decisions must take into account the conversational implicature (pragmatic effect) that each specialist would have, if used. This idea of implicature is useful to the phrasal lexicon builder: it provides him with a criterion of organization — the pseudo-syntactic class (i.e., the syntax specialist) that expresses the input *without* implicature must be the one most closely associated with it in the network.

(i.e., whether they have a verb phrase or not). Rather than follow the traditional lines of argument by constructing test sentences for each language, one can go and build a generator and note whether a number of decisions have to be made before the verb can be uttered, once the subject has been said. Certainly, for example, this is the case in English: both the sentences "she seeks the ball" and "she searches for the ball" derive from the same semantic source, yet "search" requires a preposition for the object. If a VP specialist exists, it will do the work of accessing the verb, finding the required preposition, and associating the preposition with the goal to create the object environment; after that, the verb specialist can proceed with conjugation, etc. On the other hand, if no verb phrase specialist exists, then the verb specialist will have to post the object's preposition so that, after its completion, the object specialist can find it.

Now from a linguistic point of view, there is nothing wrong with this transferral of information across specialists; however, programming experience with large systems with many interacting modules has taught that it is to be avoided². This principle of encapsulation of information is very useful to the generator builder, since it helps delimit the extent of syntax specialists. The principle, for example, allows us to conclude that English has a verb phrase. Whether or not other languages should have one is a purely practical matter: certainly this will depend on the types of decisions required to produce predicates. And, with respect to configurationality, *the notion that there exists a distinct entity called verb phrase, an entity that is the same in all languages with verb phrases, is certainly false*. It is quite conceivable that two languages each have "verb phrase" decisions that must be made across the whole predicate, though the decisions are not at all similar, resulting in different specialists.

However, not all rules of grammar deal with information that can be neatly encapsulated in parts of sentences. Some rules operate across specialists. These are the non-formative rules described earlier — the rules that deal with tense, number agreement, declension, etc., anything relating to the scoping and interrelations of syntactic environments. For example,

² "If a variable is not local it is said to be a *free variable* or a *special variable*. It is bad style to use special variables because it is difficult to understand a program if variables appear in it whose values and 'meaning' are given elsewhere." (Charnick & McDermott 85, p. 74, their italics)

number agreement (in, say, English) is scoped within sentences, and adjectival declension (in, say, German) is scoped within preposition/noun groups. That is, information about head-noun number that is determined by the subject specialist is used for conjugation by the verb specialist in English and German, and, in German, gender information that is determined by the head-noun specialist is used with article information determined by the noun group specialist to determine the appropriate endings for adjectives. This information is scoped over (has to be kept available for the duration of) the governing syntactic environment. These rules do not require separate syntax specialists. Rather, they are implicitly contained within the specialists. Thus they can apply anywhere the general and the specific formative phrases are used.

Thus, where functionally justified, syntax specialists exist, and may correspond to the traditional phrase structure entities, where each specialist creates a specific syntactic environment and, in doing so, may spawn goals to create other environments. But PAULINE's syntax specialists are not limited to the standard phrase structure symbols. The approach to organizing language presented here recognizes many more symbols than standard linguistic systems do; in fact, *any* grouping of the information and decisions that are applicable to more than a single instance may be considered a specialist. For example, English has a number of highly idiomatic ways of referring to money. This knowledge must appear somewhere in a generator's lexicon, and it seems sensible to group the indices to the relevant patterns as well as the criteria for deciding among them together in a specialist. Though not a phrase structure entity, a phrase, or a word, this specialist exists in the lexicon and performs its function similarly. PAULINE has, in its phrasal lexicon, a specialist that knows how to say (i.e., that creates a "syntactic" environment appropriate to expressing) amounts of money; another specialist for saying the time; one for colors, for measurements, etc. These specialists are described in more detail later in this chapter.

6.3.3 Relations among Elements of the Lexicon

What is the relationship between verb-based formative patterns, noun-based formative rules, and grammatical special purpose entities such as the subject specialist, and phrases such as Becker

describes? If they exist side by side in the lexicon, which ones are accessed by the generator under which circumstances? How do specific patterns differ from more general ones of the same form?

In the main, some generalities hold across the formative patterns. Very general rules, such as [SUBJECT VERB OBJECT], seem to apply at all times, even to very specific ones that express specific concepts. Sometimes, however, patterns can be ungrammatical: [*you pays your money and you takes your chance*] or [SUBJECT *ain't done nothing*]. This is consistent with the general theory of memory organization described in [Schank 82], in which special cases, exceptions, and idiosyncracies remain associated with specific episodes in memory while less specific cases are generalized to the point of maximum applicability.

Other patterns appear to be instances of general rules without in fact being so. For example, to announce Pete's demise, you can say "Pete died" or "Pete kicked the bucket", where the former uses the idiosyncratic phrase accessed directly from the state change DIE:

1. [SUBJECT [VERB *kick*] *the bucket*]

and the latter uses the general phrase

2. [SUBJECT PREDICATE]

Phrase 1 is, of course, a partially frozen expression with only one meaning. In a phrasally organized lexicon, this phrase will be indexed only under the concept DIE and nowhere else; in contrast, phrase 2 is part of the general sentence-building knowledge, able to produce many more different sentences than phrase 1 is. But that does not imply that the two phrases are different *in nature*. Both consist of an ordered list of elements, where some elements are words and others determine the syntactic environments into which words will eventually be placed. How, now, are phrases 1 and 2 related? After replacing PREDICATE in phrase 2 by its definition,

- 2a. [SUBJECT VERB OBJECT PREPOSITION-PHRASES]

phrase 1 seems to be a specialization of phrase 2. This is not, however, the case: *the bucket* in phrase 1 is not part of an OBJECT environment. Since this phrase cannot become "the red bucket" or "the buckets", *the bucket* is part of the VERB environment, making

"kick" an intransitive verb here. This explains the unacceptability of (b):

- (a) He tied the noose around his neck, kicked the chair from under himself, and kicked the bucket
- (b) * He tied the noose around his neck, kicked the chair from under himself, and the bucket

Therefore, obviously, no explicit relationship exists between phrase patterns 1 and 2.

In general, when syntactic generalities do exist between phrases and words, these generalities should be exploited. This idea was noted by Jacobs ([Jacobs 85], p 42):

...a system which deals only with "core" grammatical and productive constructs will handle but a small portion of the language... On the other hand, failing to take advantage of linguistic generalizations can introduce redundancy and possibly inefficiency into the knowledge base. Robust and efficient language processing therefore demands a balance between specialized and generalized knowledge.

As described above, this approach argues for the creation of a large number of pseudo-syntactic classes to capture the generalities. For example, Jacobs describes a concept called *transfer-event*, of which one view uses "take" and another uses "give". This transfer-event is not semantic (in PAULINE, it would be implemented as a syntax specialist). It produces the sentences

- (a) Frazier gave Ali a punch
- (b) Ali took a punch from Frazier
- (c) John gave Mary the book
- (d) Mary took the book from John

It carries semantic information, such as the fact that the punch-act and the book are the object, John and Frazier the (metaphorical) source, etc. Using the notion of *view* in the representation language KODIAK described in [Wilensky 84], Jacobs argues that transfer-event is useful for providing the necessary syntactic information to generate, in addition, sentences such as

- (e) John gave Mary a hug
- (f) John gave Mary a massage

and argues for additional restrictions to prevent

- (g) Mary took a hug from John
- (h) Mary took a massage from John

It should be clear that formative grammar rules, phrasal patterns, and syntax specialists are but slightly different incarnations of the same type of information: inclusion, casting, and ordering requirements. The differences are caused by ease of use in a system: what is a specialist function in one system with one notation may easily be a pattern in a more powerful system. In a phrasally organized system, no a priori distinction should be made between the contents of the lexicon and the contents of the grammar.

Why does one care about the relationship between the rules of grammar and the lexicon? One cares because, if all the formative grammar rules can be incorporated into the lexicon, and if the elements of the lexicon are inextricably tied to the system's network of concept representations, then the right way to build a set of representations is to pay a lot of attention to the ways in which the representation elements are expressed in language — not only the words existing for entities, but also the phrases and sentences. If the elements under consideration do not easily support such words and phrases they are suspect. Hence, generating from a representation is an excellent way of discovering its shortcomings (which is a paraphrase of McGuire's maxim: *when the generating gets tough, check that representation*).

6.4 A Phrasal Grammar

PAULINE's grammar consists of a set of phrases. Some phrases prescribe actions the generator must perform, and are implemented as specialist functions — the syntax specialists. Other phrases simply provide inclusion, ordering, and casting information and are implemented as patterns.

The phrases can be arranged in a rough hierarchy depending on how much effect they have on the final text. At the level of largest effect, the phrases control the formation of multi-predicate sentences, such as enhancer and mitigator phrases and relations between topics. At the next level, the phrases determine sentence content and organization to form various types of sentences (questions, imperatives). At lower levels, the content and organization of predicates, adverbial clauses, and noun groups are determined. Finally, words are chosen. The hierarchy does not reflect the order in which phrases actually do their work during the generation of sentences: halfway through the generation of a sentence, words will already have been chosen for some pieces of the input while other pieces are still completely uninterpreted. Furthermore, the hierarchy is not strict: for example, the choice of verb has an effect on the sentence wider than simply one word, for it often determines the presence and order of preposition groups.

The phrases that constitute PAULINE's grammar are listed in Appendix B. They can be categorized as follows:

- **Multi-predicate phrase patterns:** When appropriate, depending on the relationships between the sentence topics and on the desired slant, the topic organization phase juxtaposes topics into these patterns. Multi-predicate patterns are used to express the following:
 - **Slanting phrases** such as "Not only X, but Y" and "X, however, Y"
 - **Reminding phrases** such as "X, which reminds me of Y"
 - **Goal-relationship phrases** such as "X in order to Y" and "X so that Y"
 - **Result-relationship phrases** such as "X. As a result, Y" and "Y because X"
 - **Other relationship phrases** such as "X is larger than Y" and "After X, Y"
- **Standard phrase structure concepts** such as predicate, noun group, pronoun, etc.: Most of these require the generator to perform inclusion, ordering, and casting decisions, and are thus implemented as syntax specialist functions; the others (such as noun group form) are patterns.

- **Other standard patterns and idiomatic phrases:** Additional patterns of language such as the frozen ways of referring to money, age, titles, places of residence, greetings, expressing likes and dislikes, etc.

6.5 A Small Example of Realization

At this point, a short description of how PAULINE's syntax specialists realize a sentence will make clear the interaction of grammatical and lexical information required to produce text. The example is taken from the beginning of this chapter.

The central generator data structure is a stream -- a list of units. Of it, only the first unit (the stream *head*) is evaluated. Each unit is one of three things: a word, a topic goal, or a syntax goal. If it is a word, the unit is output and removed from the stream; if a topic goal, the topic is examined by the topic collector and the planner (as described in chapters 4 and 7) and the resulting list of syntax goals is replaced at the front of the stream. A syntax goal represents the generator's instruction to create a syntactic environment (say, a sentence, or a noun group) from the input representation element. Initial syntax goals are spawned by the planner; after that, each syntax goal produces a list of other syntax goals and/or words (this process was described as a "cascade" in [McDonald 81]), until all the words have been said. Each syntax goal contains a pointer to the specialist that produces the required syntactic environment (henceforth called the *say-function*); the representation element (henceforth called the *topic*); and other additional information relevant to the creation of the environment (pragmatic information such as desired slant, syntactic information such as gender, etc., where appropriate). Generation proceeds by applying the say-function to the topic and replacing whatever the say-function produces back on the front of the stream. This is a straightforward way to implement the left-to-right generation of language. The generator can thus be viewed as performing a depth-first traversal of the syntax tree of each sentence, where each node in the tree corresponds to a specialist function (though the tree need never be built explicitly; in PAULINE it isn't). The central expansion function simply loops until the stream is empty.

During expansion, the generator must maintain certain information for certain periods of time. This information helps satisfy the non-formative requirements of the grammar. For example, for the duration of a noun group, the generator must have available information about case (and in languages with more explicit declension such as German and French, about number and gender); for the duration of a sentence (a "verb group"), it requires information on number and mode (for subject-verb agreement), tense (for verb-time clause agreement), etc. The requisite information for noun groups and sentences is maintained in *context records* on two stacks; when a noun group is started, a new context is created on the noun group context stack and given values (case, etc.) by the relevant specialist functions; when the noun group is finished, the dummy say-function POP-NOUN-GROUP pops the context. For pronominalization, the generator also creates a list of the topics it has already said, together with appropriate syntactic information such as gender and case.

6.5.1 Input Topic

We begin with one syntax goal on the stream:

```
# {SYNTAX-GOAL
  [SAY-FUNCTION : SAY-SENT-TOP]
  [TOPIC : MTRANS-6]}
```

where MTRANS-6 is the following (JIM, JANET, QUIET, HEALTH, etc., are all either atomic symbols or defined in the system's representation network):

```

#{ACTION = MTRANS-6
  [ACT : MTRANS]
  [ACTOR : JIM]
  [OBJECT : #{STATE-CHANGE = DEATH-10
    [TYPE : HEALTH]
    [ACTOR : JANET]
    [FROM : ALIVE]
    [TO : DEAD]]}
  [FROM : JIM]
  [TO : #{PERSON = SUE
    [AGE : #{MEASURE = AGE-23
      [UNIT : YEAR]
      [NUMBER : 23]]}
    [NAME : SUE]
    [RESIDENCE : NEW HAVEN]
    [SEX : FEMALE]
    [SIZE : SMALL]]}
  [MANNER : QUIET]}

```

6.5.2 Sentence Content and Organization

The generator's expansion loop always expands the stream head. It starts by applying SAY-SENT-TOP to MTRANS-6. This specialist function must determine what type of sentence to make; since MTRANS is an event, it returns the following syntax goal to be placed on the stream:

```

#{SYNTAX-GOAL
  [SAY-FUNCTION : SAY-EVENT-SENT]
  [TOPIC : MTRANS-6]}

```

Since this goal is the head entity on the stream, it is expanded next. Its specialist SAY-EVENT-SENT is applied to MTRANS-6. First, it checks whether the rhetorical strategies call for nominalization of the input topic; if so, it returns

```
# {SYNTAX-GOAL
  [SAY-FUNCTION : SAY-PRONOUN]
  [TOPIC : MTRANS-6]}
```

which would eventually expand to "that", as in "That was what happened". Otherwise, the specialist performs the following tasks: it selects a subject, selects pre-subject adverbial aspects, and sets the sentence context information for tense (from the TIME aspect, if any) and mode (PASSIVE when the subject is the OBJECT (or, in an MTRANS, the TO) aspect; ACTIVE otherwise). The criteria for subject selection are both syntactic (using rules of topic coherence such as those discussed in chapter 4) and pragmatic (such as the rules described under **RG:formality** in chapter 5). The criteria for including pre-subject clauses and selecting from the available candidates are also pragmatic. In the example, only one adverb, MANNER, has been given; other possibilities are TIME, LOCATION, and INSTRUMENT. If the pragmatic criteria call for its inclusion, and decide to make JIM the subject, this specialist returns:

```
# {SYNTAX-GOAL
  [SAY-FUNCTION : SAY-PRE-SENT]
  [TOPIC : MTRANS-6]
  [ASPECTS : (MANNER)]}

# {SYNTAX-GOAL
  [SAY-FUNCTION : SAY-SUBJECT]
  [TOPIC : JIM]}

# {SYNTAX-GOAL
  [SAY-FUNCTION : SAY-PREDICATE]
  [TOPIC : MTRANS-6]}
```

This sequence of syntax goals will eventually expand into sentences such as "Quietly, Jim told Sue...". Alternatively, under appropriate pragmatic circumstances, the rhetorical

criteria prescribe the selection of DEATH-10 as the sentence subject, which would eventually produce "Janet's death was what Jim told Sue of". Or, if the criteria suppress all pre-subject adverbials and select SUE as the subject, the result is:

```
# {SYNTAX-GOAL
    [SAY-FUNCTION : SAY-SUBJECT]
    [TOPIC : SUE]}
# {SYNTAX-GOAL
    [SAY-FUNCTION : SAY-PREDICATE]
    [TOPIC : MTRANS-6]}
```

6.5.3 Clause Organization

The specialist SAY-SUBJECT creates a new noun group context in the nominative case and pushes it on the noun group context stack. It checks whether the topic is a single entity or whether it should build a list of entities (as in "Thomas, Richard, and Harold went..."), and returns one or more goals:

```
# {SYNTAX-GOAL
    [SAY-FUNCTION : SAY-NOUN-GROUP]
    [TOPIC : SUE]}
# {SYNTAX-GOAL
    [SAY-FUNCTION : POP-NOUN-GROUP]}
```

SAY-NOUN-GROUP then queries the rhetorical strategies to select a head noun from the aspects PAULINE's grammar can handle (namely, TYPE, OCCUPATION, TITLE, NATION, RESIDENCE, AGE, NUMBER, GENDER, LOCATION, and SIZE, in that order, all other things being equal), as well as from the generator's own opinions (if any), using syntactic and pragmatic criteria, as described in chapters 4 and 5. This specialist also determines pre- and post-nominal modifiers. In this regard, some possibilities are ruled out by syntactic text flow rules (one doesn't say "the New Haven female 23-year-old Sue

Brown, a comptroller"); other decisions are made by strategies that query the activated rhetorical goals and the affective values of the aspects. These strategies determine the affective and knowledge-related suitability for each of the candidate modifiers, as described immediately below, and return the permissible candidates in order of preference. Then SAY-NOUN-GROUP selects some number of these candidates, depending on its syntactic adjective combination rules and the value of the rhetorical goals **RG:detail** and **RG:haste** (for example, when the latter goal is *pressured*, no aspects are included, as described in chapter 5). Finally, the specialist associates each selected candidate with an appropriate say-function specialist (maintaining the preferred order), syntactic ordering constraints are checked, and the resulting list of syntax goals is returned.

The pragmatic topic inclusion strategy determines the knowledge-related and affective suitability of a given input element as follows: First, if the element is directly inferable from the current topic, then saying it is not very important. This decision is, of course, very complex; for relevant work see [Appelt 82, 87] and [Cohen 78]. PAULINE simply checks whether the element is directly related to the topic; if so, it checks whether it is knowledge that the hearer is known to have, and whether conversational conditions are not bad (noisy) enough to warrant saying the element nonetheless. Second, if the input element inherits an affect from the interlocutors' sympathies, then the rhetorical goal **RG:partiality** is checked. If the program's own affect for the element is not the same as its intrinsic affect (as defined in chapter 4), the importance of its being said depends on whether **RG:timidity** is *reckless* or whether **RG:respect** is *arrogant*. If the hearer has the same affect for the input element as the speaker does, the strategy checks **RG:aggression** and gives the element high importance if this goal requires friendship (i.e., if it has the value *placating*).

Different types of topic require different noun group forms: locations ("X was where..."); states ("Peter's feeling ill..."); possessives ("the man's large car..."). This specialist also has pragmatic strategies select and order pre- and post-nominal modifiers. For example, SAY-NOUN-GROUP could produce:

```

#{SYNTAX-GOAL
  [SAY-FUNCTION : SAY-ARTICLE]
  [TOPIC : SUE]}
#{SYNTAX-GOAL
  [SAY-FUNCTION : SAY-PRE-NOUN-MODS]
  [TOPIC : SUE]
  [ASPECTS : (SIZE)]}
#{SYNTAX-GOAL
  [SAY-FUNCTION : SAY-HEAD-NOUN]
  [TOPIC : SUE]
  [ASPECT : AGE]}
#{SYNTAX-GOAL
  [SAY-FUNCTION : SAY-POST-NOUN-MODS]
  [TOPIC : SUE]
  [ASPECTS : (RESIDENCE NAME)]}

```

6.5.4 Word Choice

At this point, then, the expansion stream contains:

```

#{SYNTAX-GOAL
  [SAY-FUNCTION : SAY-ARTICLE]
  [TOPIC : SUE]}
#{SYNTAX-GOAL
  [SAY-FUNCTION : SAY-PRE-NOUN-MODS]
  [TOPIC : SUE]
  [ASPECTS : (SIZE)]}
#{SYNTAX-GOAL
  [SAY-FUNCTION : SAY-HEAD-NOUN]
  [TOPIC : SUE]
  [ASPECT : AGE]}

```

```

#{SYNTAX-GOAL
  [SAY-FUNCTION : SAY-POST-NOUN-MODS]
  [TOPIC : SUE]
  [ASPECTS : (RESIDENCE NAME)]}
#{SYNTAX-GOAL
  [SAY-FUNCTION : POP-NOUN-GROUP]}
#{SYNTAX-GOAL
  [SAY-FUNCTION : SAY-PREDICATE]
  [TOPIC : MTRANS-6]}

```

The definite article is used for nouns such as "police", for topics that already have been said, for topics that have enough specifying information (such as specific time or place), or for phrases such as "the most". The article is suppressed when the topic is commonly known or is a proper noun (such as "Mexico"). "That" is used with an explicitly opinionated noun ("that jerk"). Otherwise, as in the example, the specialist SAY-ARTICLE returns the word

====> THE

This word is popped off the stream and said. The next specialist, SAY-PRE-NOUN-MODS, casts the aspect into an appropriate say-function:

```

#{SYNTAX-GOAL
  [SAY-FUNCTION : SAY-SIZE]
  [TOPIC : SMALL]
  [NG-POSITION : PRE]}

```

which produces

====> SMALL

The SAY-HEAD-NOUN specialist finds an appropriate word or phrase from the representation element (perhaps using pragmatic factors in its discrimination), and then pluralizes if necessary. Sometimes, as in this example, it requires another specialist:

```
# {SYNTAX-GOAL
  [SAY-FUNCTION : SAY-AGE]
  [TOPIC : AGE-23]
  [NG-POSITION : HEAD]}
```

```
====> 23-YEAR-OLD
```

SAY-POST-NOUN-MODS casts adjectival aspects using appropriate specialists:

```
# {SYNTAX-GOAL
  [SAY-FUNCTION : SAY-RESIDENCE]
  [TOPIC : NEW HAVEN]
  [NG-POSITION : POST]}
```

```
# {SYNTAX-GOAL
  [SAY-FUNCTION : SAY-NAME]
  [TOPIC : SUE]
  [NG-POSITION : POST]}
```

which produce

```
====> FROM NEW HAVEN
====> *CMA* SUE *CMA*
```

Now the dummy specialist POP-NOUN-GROUP is encountered and the noun group context stack is cleared.

The next task is to build the predicate. The most important task here is to select a verb; as described before, all the formative information can then be found -- either from it, or from the representation element, or by using the default form. To find verbs in the lexicon, PAULINE uses discrimination nets attached to its representation primitives. Its searches the concept definition network 'near' the topic, starting with the type of the topic and proceeding up the property inheritance hierarchy until a filled WORD aspect is found; this will point to a verb or to a (discrimination) procedure that will eventually point to a word.

Near the top of the hierarchy all elements are organized as some Conceptual Dependency primitive.

PAULINE's lexicon contains more than 20 words to express MTRANS; amongst others, "tell", "inform", "whisper", "shout", and "broadcast". Since the mode in this example is passive, some verbs (such as "say") are not available and others (such as "hear"), normally unavailable, are. The discrimination criteria are semantic and pragmatic, as described earlier and in chapters 4 and 5. If the verb "hear" is selected, the generator must use the sentence form

```
[ VERB [PREPGROUP (aspect OBJECT) of]
  [PREPGROUP (aspect ACTOR) from] SAY-POST-SENT ]
```

whereas the verb "tell" specifies the form

```
[ SAY-VERB [SAY-OBJECT (aspect TO)] SAY-ADVERB SAY-POST-SENT
  that [SAY-COMPL (aspect OBJECT)] ]
```

Choosing the latter, and inverting for the passive mode, this specialist returns

```
#|SYNTAX-GOAL
  [SAY-FUNCTION : SAY-VERB]
  [TOPIC : MTRANS-6]
  [WORD : tell]
  [MODE : PASSIVE]}
#|SYNTAX-GOAL
  [SAY-FUNCTION : SAY-PREPGROUP]
  [TOPIC : JIM]
  [PREPOSITION : by]}
that
#|SYNTAX-GOAL
  [SAY-FUNCTION : SAY-COMPL]
  [TOPIC : DEATH-10]}}
```

The specialist SAY-VERB conjugates the verb, producing

```
====> WAS TOLD
```

and SAY-PREPGROUP expands via SAY-NOUN-GROUP directly into SAY-HEAD-NOUN:

```
====> BY JIM
```

The next element on the stream is the literal "that", which is just said:

```
====> THAT
```

6.5.5 The Process Repeats

Only one syntax goal remains: the clause expressing DEATH-10. Its syntax specialist, SAY-COMPL, creates a sentence without initial adverbial clauses, linking words, and opinions, by activating SAY-SENT-TOP with appropriate switches (which are simply listed here as part of the syntax goal):

```
#|SYNTAX-GOAL
  [SAY-FUNCTION : SAY-SENT-TOP]
  [TOPIC : DEATH-10]
  [SUPPRESS : (PRE-ADVERBIALS LINK-WORDS OPINIONS)]
```

As before, SAY-SENT-TOP must determine what kind of sentence to make. A state-change is made by using SAY-EVENT-SENT, which expands into

```
#|SYNTAX-GOAL
  [SAY-FUNCTION : SAY-SUBJECT]
  [TOPIC : JANET]}
#|SYNTAX-GOAL
  [SAY-FUNCTION : SAY-PREDICATE]
  [TOPIC : DEATH-10]}
```

Proceeding as before, SAY-SUBJECT and SAY-PREDICATE expand into

```
====> JANET
====> DIED
```

Alternatively, since the frozen phrase "kick the bucket" is associated with the state-change DEATH, the generator can select this option instead (under appropriate pragmatic conditions such as, say, *colloquial* **RG:formality**):

```
# {SYNTAX-GOAL
  [SAY-FUNCTION : SAY-VERB]
  [TOPIC : DEATH-10]
  [WORD : kick]
  [MODE : ACTIVE]

  the
  bucket
```

thereby expressing its irreverence.

6.6 Conclusion

In this chapter, we ask: How should language be represented in a generator program? In particular, how do the concepts the generator must express, the grammar it is to use, and the words and phrases with which it must express them, relate? The answer presented here is that all linguistic knowledge — all language — should be contained in the lexicon. The argument is the following: The generator must perform three types of task to produce text. It gets the information it requires to do these tasks from three sources: from the grammar, from partially frozen phrases (including multi-predicate phrasal patterns), and from certain words.

Verbs, nouns, and other words in the lexicon are associated with the representation elements they describe. When idiosyncratic forms of expression exist, the formative patterns

are associated with the lexical entries. Similarly, frozen and semi-frozen phrases describe representation elements, are associated with representation elements in the same way, and have similar formative patterns. Formative patterns are built out of other patterns and lexicon entries. Although the general formative rules of grammar are not associated with any specific representational element, their formative patterns are defined similarly. From the generator's perspective, there is no difference between these sources: for example, there is no real reason for the grammar of the language to be contained in a separate "syntax box". Thus rules of grammar, multi-predicate phrases, and phrasal and verb predicate patterns can all be viewed as phrases, frozen to a greater or lesser degree, and should all be part of the lexicon. Some such "phrases" can be quite complex, prescribing a series of actions and tests: these can be thought of as specialist procedures. Others can be very simple: templates. These specialist elements contain the information that enables PAULINE to perform the three tasks of inclusion, ordering, and casting, in order to expand the input into a string of words. The elements that constitute PAULINE's lexicon, and the way they are are used, are described.

Chapter 7

Planning and Realization

Abstract

Traditional (blocks-world) hierarchical expansion planning is not suitable for all planning tasks in language generation. A more appropriate method, limited-commitment planning, consists of both prescriptive (blocks-world) planning and of restrictive planning (selecting from options with reference to the status of active goals); it is interleaved with the realization process. At present, existing text planners use prescriptive plans exclusively. However, a large class of planner tasks, especially those concerned with the pragmatic content of text such as style and slant, is most easily performed under restrictive planning. The kinds of tasks suited to each planning style, and the way PAULINE performs limited-commitment planning, are described.

7.1 The Trouble with Traditional Planning


In the traditional planning paradigm, one (or more) initial goal(s) is transformed, after a hierarchical goal-plan expansion cycle, into a series of steps that are executed by some agent (see, say, [Sacerdoti 77]). In the generation process, this approach takes the form of building up and associating generator instructions with increasingly detailed plans.


AD-A183 418 GENERATING NATURAL LANGUAGE UNDER PRAGMATIC CONSTRAINTS 3/3
(U) YALE UNIV NEW HAVEN CT DEPT OF COMPUTER SCIENCE
E H HOVY MAR 87 VALEU/CSD/RR-521 N00014-85-K-0108

UNCLASSIFIED

F/G 5/7

NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

input topics until, eventually, enough instructions have been assembled to realize each part of the input as one or more words. A number of planners base their decisions mainly on the state of the hearer's knowledge: [Appelt 81], [Woolf & McDonald 84], [McCoy 87]. Some planners are concerned with the selection and coherent presentation of specific information: [McKeown 82], [Paris & McKeown 87], [Rösner 87], [Mann 84]. A few planners seek to achieve various other pragmatic effects: [Cohen 78], [Jameson 87], [Bienkowski 86].

Of course, planning all the way down to the actual details of word choice requires that the planner have access to as much syntactic knowledge as the realization component itself. This obviates the need for a realization component. For example, suppose the generator wants to create in the hearer sympathy for a 65-year old beggar. In the sentence "the [SAY-AGE #AGE-1] woman is homeless", the specialist SAY-AGE should return "old" or even "ancient" rather than "65-year old". For the planner to precompute this decision, it will have to compute all the decisions (via SAY-SENTENCE and SAY-SUBJECT, etc.), such as selecting a subject, a head noun, and adjectives, before it will be in a position first to realize that #AGE-1 is to be said as an adjective, and second to determine what the options are in this case. In order to do this computation, the planner will have to have access to information which one would like to claim is properly the exclusive concern of expansion, such as syntactic and lexical knowledge. (For instance, Appelt's planner contains grammatical knowledge spread throughout. Appelt alludes to the problems that this causes in [Appelt 81, p 113].) If the planner is going to do all this work, down to the level of individual words, it may as well do the generation simultaneously. This model is unwieldy: it mixes planning and realization information.

Furthermore, this model is unrealistic: When we speak, we do not try to satisfy only one or two goals, and we operate (often, and with success) under conflicting goals for which no resolution exists. We usually begin to speak before we have planned out the full utterance, and then proceed while performing certain planning tasks in bottom-up fashion. That is, when we start speaking, we have usually made some decisions and have postponed others — we have some vague notion about what topics we want to cover, and maybe even of the desired slant and a particular phrase we want to use; we leave the details — especially the syntactic details — for later, real-time, consideration.

A better solution is to perform planning only when necessitated by the expansion. This approach is characterized by a two-way communication at decision points. In other words, this suggests that the planner assemble only a partial set of generator instructions — enough for the realization component to start working on — and then continue planning when the realization component requires further guidance. Thus, as argued in [Hovy 85] and in [McDonald & Pustejovsky 85], the solution is to interleave planning and realization. In addition, a number of psycholinguistic experiments that investigate the components of human generation processes also conclude that the model of strictly sequential separate components (planning, realization) is incorrect; see, say, [Rosenberg 77] and [Danks 77].

This is a fine solution; however, there still is a problem. It is impossible to formulate workable plans for certain types of goals that speakers frequently have. This is true especially for pragmatic goals. A speaker may, for example, have the goals to impress the hearer, to make the hearer feel socially subordinate, and yet to be relatively informal. These goals play as large a role in generation as the speaker's goal to inform the hearer about the topic. However, they cannot be achieved by constructing and following a plan — what would the plan's steps prescribe? Certainly, making explicit sentences such as "I want to impress you, but still make you feel subordinate" would be counterproductive. Pragmatic effects are best achieved by making subtle, small decisions during the generation process: an extra adjective here, a slanted verb there. But how can a generator control this?

The problem stems from the fact that planning is usually prescriptive: it determines multiple actions over a long range of time (or text). However, when a system is busy executing a partially formed plan, and the only decisions left to be made are very limited in range, the nature of the task changes: the system is faced with a series of unrelated choices from sets of options. This task is usually called "selection" rather than "planning". As such, it may not seem very important, as long as the planner has assembled detailed enough instructions for each selection to be made. For example, in the blocks world, the actual realization of a low-level goal such as [MOVE HAND TO BLOCK A] may not matter (that is, the hand may equally well move along any reasonable path). However, in generation, low-level realization decisions can have significant pragmatic effects; for example, compare saying "terrorist" to "guerrilla" or "freedom fighter". Hence the problem: the planner

cannot assemble detailed enough instructions for each pragmatic choice, because it cannot know beforehand which choices it will face — remember, it doesn't have access to syntactic information — and thus which pragmatic goals are likely to be affected by each selection. Also, it cannot simply plan for all the goals together, because the order of satisfaction of goals may be important. Unlike hierarchical plan steps, adjacent selection decisions need not work toward the same goal (or indeed have any relation with each other); the planner has no way to guess even remotely what the next selection or satisfiable goal might be. Thus hierarchical plan expansion is not appropriate for this selection-type "planning". What is required is something quite different: strategies that guide selection decisions, based on the state of satisfaction of pending goals.

7.2 Limited-Commitment Planning

Generation, then, requires two types of planning. Certain tasks are most easily performed in top-down fashion (that is, under guidance of a hierarchical planner, or of a fixed-plan (schema or script) applier), and other tasks are most naturally performed in a bottom-up, selective, fashion. That is, some tasks are *prescriptive* — they act over and give shape to long ranges of text — and some are *restrictive* — they act over short ranges of text, usually as a selection from some number of alternatives. Our generators' planners should not be solely prescriptive; they should also handle restrictive information, in the sense of considering information about options that is not contained in any explicit plan but that is brought up by the (state of the) realization process.

The difference between prescriptive and restrictive planning is captured neatly in the location, organization, and use of the information required to guide the generation process from its initial goal (to say the topic) to a final state (the said utterance). When the information is contained in a set of plan steps (ordered or unordered), and when activated this information directs which specific actions are done, then the plan is prescriptive. On the other hand, when the information is contained elsewhere (in, say, the lexicon), and the planner uses this information to query its goals for guidance, then this information

(together with the strategies used for resolving goal conflicts, etc.), constitutes restrictive plans. Prescriptive information is formative: it controls the construction and placement of parts in the paragraph and the sentence. Under guidance of its prescriptive plans, the generator makes some commitment to the final form of the text (such as, for example, the inclusion and order of specific sentence topics). However, it leaves open a number of decisions (such as, for example, the possibility of including additional topics under certain conditions, or the specific content of each sentence), until actual realization takes place. These unrealized decisions are handled by restrictive plans, which are selective: they decide among alternatives to build the actual text.

This planning paradigm can be called *limited-commitment planning*. Limited-commitment planning applies not only to generation; it is ubiquitous in daily life. We very seldom do pure hierarchical expansion planning. For example, the limitations of pure top-down planning in argumentation are described in [Birnbaum 86, p 176]:

An exclusively top-down approach to planning can work in situations which are more or less under the control of the planner... But conversations do not, in general, meet those requirements... Thus, unless a speaker can predict, rather specifically, how his adversary will respond, his utterances cannot be completely planned in advance.

McDermott describes a general problem solver/planner that uses a theorem prover to retrieve plan schemata for attacking problems. When it finds more than one schema, it retrieves choice rules to help select one. Here "task reductions" correspond to prescriptive plans and "choice rules" to restrictive plans [McDermott 78, p 76]:

... sometimes the user will want to be able to express rules for synthesizing a brand new alternative task reduction 'on the fly' when two task reductions have been suggested... The solution is to face up to the necessity for treating 'choice between alternatives' as a basic situation for problem solving.

7.3 Performing Restrictive Planning: Monitoring

Due to its bottom-up, run-time nature, restrictive planning differs from traditional planning in the type of information it requires and in the way it uses that information. For example, limited-commitment planners sometimes have to face up to and compromise on goal conflicts, rather than employ procedures called *critics* (in NOAH, [Sacerdoti 77]) to notice when conflicts occur and to plan around them. For example, a conflict can arise when the speaker instructs someone who is socially dominant: since he is instructing, he has stylistic permission to change the topic whenever he thinks it appropriate; but since he is subordinate to the hearer, the hearer should control topic change. For traditional planners such as NOAH this conflict poses a serious problem. However, this is what gives generation its spice! People can speak while holding conflicting goals; generators must be able to merge conflicting plans into sensible generator instructions.

Thus, an important difference between the two types of planning is that a restrictive planner is unable to guarantee that all its goals will necessarily be achieved; it can, by making appropriate decisions, merely ensure that no goal gets too seriously thwarted. That is, it cannot simply plan *for*, it is constrained to plan *with*: the options it has to select from are *presented to it* by some other component (in generation, the realizer). It cannot even use current options to control future decisions, since the options only have relatively local effect, and since in addition the information that controls the sequence of decision types and hence of the options (in generation, grammar), is not available for planner inspection. Thus a restrictive planner deals with a number of goals whose satisfaction statuses are constantly changing in ways it cannot predict. What information, then, does it require?

Obviously, after each decision, the statuses of the affected goals must be altered. This task is called execution monitoring in a real-time planning system with an agent (see, say, [Fikes, Hart & Nilsson 72], [Sacerdoti 77], [Miller 85], [Doyle, Atkinson & Doshi 86]); we will use the term *monitoring* here, appropriate for a system that does not take into account the world's actual reaction (in generation, the hearer's actual response), but that trusts, perhaps naively, that the world will react in the way it expects. Monitoring can be

performed by the restrictive planner, since it deals with the options' effects on the goals. This task requires checking, updating, and recording

- 1. the current satisfaction status of each goal

Furthermore, to perform both planning and monitoring, the restrictive planner must know (or be able to compute), for each option,

- 2. which goal(s) it will help satisfy, to what extent, and in what ways
- 3. which goal(s) it will thwart, to what extent, and in what ways

— that is, it must be able to judge the effect of each option as far as permitted by restrictive planning. In practice, of course, this cannot be too far, because if the planner tries to increase the extent by replacing the information it is given with more detailed information, from whatever source, it will be performing top-down (i.e., prescriptive) planning. (Thus, in generation, the way additional information is used is the principal line of demarcation; other than this, there is no fixed point, no natural division in the generation process before which a prescriptive planner should run and after which a realizer (and with that, a restrictive planner) should take over). Hence the restrictive planner must be provided not only with options, but also with some indication of which goal(s) each option will affect and in what way it will do so. Obviously, this information should reside in the subsystems or specialist functions that present the options to the planner.

Finally, in order to perform restrictive planning, the planner must be able to resolve conflicts when various options satisfy various goals; it therefore requires

- 4. the relative priority of each goal

As a partial solution, [Durfee & Lesser 86] present **three heuristics** to help a planner that is uncertain about which long-term goals to pursue and which sequence of actions to select. These heuristics are: prefer common intermediate goals; prefer cheaper goals; and prefer discriminative intermediate goals (goals that most effectively indicate the long-term promise of the avenue being explored). In addition, two simple strategies for determining

priority are **least-satisfied** and **least-recently satisfied**. A better alternative is the **combination of these two** strategies, in which a goal receives higher priority the longer it waits and the fewer times it has been satisfied. The introduction of explicit priorities permits the formulation of additional high-level prescriptive plans to control the planner's overall approach. (In generation, this is expressed as the *tenor* of longer texts; for example, in a speech, the speaker may start off distant and forceful, and then gradually become more friendly and relaxed. Appropriate plans, when worked out, may be implemented in the style of the meta-plans described in [Woolf & McDonald 84].)

Restrictive planning is implemented in PAULINE in the following way: None of the program's pragmatic goals are ever fully achieved and flushed; they require decisions to be made in their favor throughout the text. (That is to say, you cannot at some point simply flush the goal to be appropriately formal, but you can give it a lower priority for a while.) An option can affect a goal in one way only — by adding one "point" to its satisfaction status — so that the satisfaction level of each goal is simply the number of times some option that helps achieve it has been selected. Of course, a single option may help satisfy a number of goals concurrently. When making decisions, the program compares the effect of each option — the total of all increased satisfaction statuses — and chooses the option corresponding to the goals with the lowest total. In order to do this, it must know which goals each option will help satisfy. As discussed above, responsibility for providing this information lies with whatever produces the option: either the lexicon or the language specialist functions in the realizer (such as a predicate or clause constructor). Thus PAULINE uses *least-satisfied* as a conflict resolution strategy. This implementation rests on the assumptions that all the program's rhetorical goals are equally important and that every option helps satisfy its goals by the same fixed amount. After every decision, the program increases the satisfaction status of each affected rhetorical goal. These simplifications produce a quite reasonable approximation of the stylistic behavior of speakers; that is, analysis of spoken texts indicates that people often alternate the strategies when pursuing conflicting goals.

By its bottom-up nature, restrictive planning provides the generator with a kind of opportunism (a very limited version of the kind described as slips of the tongue in [Freud 35] and discussed in [Birnbaum 86]). Whenever the restrictive planner selects options

that help satisfy a number of goals and that, in addition, uncover new possibilities (such as for additional topic inclusion), it can be said to be acting opportunistically. For true opportunism, a system has to check incoming possibilities and their effects against *all* its goals; as Birnbaum points out, a potentially expensive operation. However, when each option is explicitly marked with the goals it can help satisfy, the resulting limited form of opportunism becomes tractable.


7.4 Planning in PAULINE

The rest of this chapter describes the architecture and implementation of PAULINE's planning and realization processes. Each of the planning phases — additional topic collection, interpretation, etc. — is discussed, with specific reference to the prescriptive or restrictive character of the plans used. An annotated trace of a session with PAULINE appears in appendix A.

7.4.1 Program Architecture

The program consists of over 12,000 lines of T, a Scheme-like LISP developed at Yale. As input, before each run, PAULINE is set up with descriptions of the situation and of the hearer, as well as with appropriate pragmatic goals. It is also given the principal topics of conversation. Its structure is:

Program Architecture



7.4.2 Computing Affect

As described in chapter 4, opinions are modelled in PAULINE using three affects: GOOD, NEUTRAL, and BAD. The program gets its affects from two sources. The first source is the user: to give PAULINE opinions, the user must list one or more representation elements as *sympathies* or as *antipathies*. The second source is the default intrinsic affects associated with concepts; affects must be defined for certain representation elements. In the JUDGE

examples, when PAULINE defends Mike, the sympathy list contains the concept "Mike" and the antipathy list the concept "Jim". In the Carter-Kennedy examples, PAULINE is made a Carter supporter by marking the element representing Carter's goal to win the nomination as GOOD and marking Kennedy's goal as BAD. In the shantytown examples, PAULINE is given one protagonist and his goals as sympathies and the other and his goals as antipathies.

In order to determine its opinion about any arbitrary piece of input representation, the given affects must be combined with the concepts' intrinsic affects and must be propagated along the relations between concepts. PAULINE's rules are:

- affect is preserved when combined with NEUTRAL
- like affects combine to GOOD
- unlike affects combine to BAD
- *when the two affect-bearing concepts are related in certain ways the combined affect inverts (for example, when one affect bearer is a BAD action and the other is its conceptual patient). Rules of affect propagation are defined for each possible relationship between affect-bearing concepts*

This works as follows: assume the current topic is the action

#{ACTION-UNIT = ACT-6

[ACT : HIT]

[ACTOR : MIKE]

[OBJECT : BODY-PART]

[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]]}]

```

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

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

In order to slant this, PAULINE has to determine the affect for Mike's role in ACT-6 from its sympathies. If it likes Mike, the program has to combine 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 by the affect rule (discussed immediately below) the action should be omitted (not possible in this example, since there is only one action) or should be mitigated. One of the ways PAULINE could say this is:

(b) MIKE JUST TAPPED JIM ONCE

If, on the other hand, PAULINE's sympathies are for Jim, then 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, in ACT-6, Jim looks good (ACT-6 is GOOD for the case against Mike). In this case the affect rule suggests that PAULINE 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). The 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 the sentence

(c) MIKE PURPOSELY SMASHED JIM AND KNOCKED HIM DOWN

Affect propagation works similarly in the shantytown and Carter-Kennedy domains. For example, the following represents the disassembly of the shantytown:

```
#{CONSTRUCT = CONSTRUCT-2
  [ACTOR : OFFICIALS]
  [OBJECT : SHANTYTOWN-1]
  [MODE : DISASSEMBLE]
  [LOC : BEINECKE-PLAZA-1]
  [TIME : TIME-5]
  [RELATIONS :
    (REL-SUBGOAL-TO :
      #{HAVE-GOAL = GOAL-2
        [ACTOR : YALE]
        [DESIRE : STATE-ORDERLINESS]
        [TIME : PAST]
        [RELATIONS : ... ]}
    REL-WHILE :
      #{M-ALTER-FREEDOM = ARREST-1
        [ACTOR : POLICE]
        [OBJECT : STUDENTS-2]
        [MODE : ARREST]
        [TIME : TIME-5]
        [RELATIONS : ... ]}
    REL-RESULT :
      #{SUPPORT-2 ... }
```

REL-RESULT :

#{MTRANS-7 ... }]]}

When PAULINE speaks as a protester, its sympathy list contains STUDENTS-1 (the protesters) and GOAL-1 (the protesters' goal that Yale divest), and its antipathy list contains YALE and Yale's goal GOAL-2 (Yale's goal that the university remain in an orderly state). To work out its affect for CONSTRUCT-2, the program first negates the intrinsic affect of a CONSTRUCT (which is GOOD) because the action is a disassembly, and then combines this BAD affect with its affect for the actor, which is NEUTRAL. Finally, it combines the result, BAD, with the affect of the goal that the action serves, which is BAD, to get the result GOOD — that is to say, CONSTRUCT-2 should be included, and even stressed, in the text. In contrast, when PAULINE speaks as a Yale official, the resulting affect for CONSTRUCT-2 is BAD. So is ARREST-1, the arrest of the 76 students, after similar derivation. Both topics should be mitigated or avoided altogether, if possible.

7.4.3 Rhetorical Goals of Opinion

PAULINE's rhetorical goals of opinion are listed in chapter 2; these goals determine the use of strategies to make one side look good by, for example, describing how they help other people; how their side does good actions to achieve their goals; how they react to the opponents' actions in a reasonable way; and how other people support them. In contrast, strategies to make the opponents look bad are to show how their side has unacceptable goals; to explain how their side does bad actions to achieve their goals; to show their unreasonable reactions to overtures for a settlement; and to show how nobody likes them.

In more detail, some of these subgoals are defined as follows (the terms in capitals are elements of the representation language, which is based on Conceptual Dependency Theory (see [Schank 72] and [Schank & Abelson 77]): MTRANS denotes the act of transferring information; PTRANS the act of transferring physical objects, and ATRANS the act of transferring control over something):

The opponents are unreasonable because

- *Coercion*: they force their will on others (corresponding to the university speaker's "wanted to force"):

IF the current topic is an ACTION,
 AND it is an antipathy,
 AND the action serves one of the opponents' goals
 AND the goal is to have some other party do some act
 AND the other party is a sympathy
 THEN imply that the opponents force their will on them (using
 verbs and phrases such as "force", "make them do")

- *Limited support*: they claim to have more support than they have (corresponding to the university speaker's adjective "a small group"):

IF the current topic claims support (an MTRANS of a SUPPORT),
 AND the ACTOR is an antipathy,
 AND the SUPPORT contains a number of people,
 THEN minimize that number by using minimizing adjectives
 such as "a small number", "a few"

- *Appropriation*: they use distasteful/ugly tactics, misuse their rights, or overstep the bounds of propriety (which causes the university speaker to say "force" and "take over"):

IF the current topic is an ACTION,
 AND it is an antipathy,
 AND the ACTOR is an antipathy,
 THEN imply that the action is ugly, by
 IF an ATRANS or PTRANS of props for other actions:

say that they take control of what they need by force
 (use appropriate nouns, verbs, and phrases such as
 'take over', 'grab', 'commandeer')

IF an MTRANS of demands: say that they try to coerce others
 (use verbs such as 'force', 'pressure')

In contrast, we are reasonable because

- *Good goals*: we are looking after everyone's interests, not just our own; we provide a service for the community; our goals are good

IF the current topic is an ACTION or has a desired STATE
 AND the ACTION or desired STATE serves a GOAL
 AND the beneficiary of this goal is some other group
 THEN say the topic, the goal, and the beneficiaries

- *Support*: we have a lot of support; most people agree with us

IF the current topic is an MTRANS
 AND we are the recipients (it is TO us)
 AND the message (its OBJECT) expresses support
 THEN state the support, mentioning the senders

also,

IF some event or action served one of our goals
 AND some number of people participated as ACTORS
 THEN mention the event, maximizing the number of participants

- *Sadness*: the opponent's action saddened us

IF the opponents are the ACTORS of the current topic
 AND the topic serves some GOAL
 AND this goal opposes some GOAL we have
 AND we did not retaliate immediately
 THEN say outright that the action disappointed us

These goals affect the generation process in a wide variety of ways. Correspondingly, they are best implemented in various ways. Some goals are more naturally implemented top-down: for example, the goals that prescribe the inclusion of sentence topics that are not necessarily part of the input topics at all, such as the sentences "we deplore the whole affair" and "they started it". These sentences are added purely for slanting effect; it is doubtful that any non-generator process (other than an argument builder) would routinely include such topics in its output. Other goals function best in bottom-up fashion: They respond to syntactic opportunities to include aspects; for example, the goals that suggest the inclusion of appropriate stress words and the selection of appropriate adjectives, nouns, and verbs (as described in chapter 4) cause the production of such phrases as "a few shiftless students".

7.4.4 Topic Collection

The task of collecting appropriate additional sentence topics is pre-eminently prescriptive. The best examples of plans for this task are the schemas developed by McKeown ([McKeown 82], [Paris & McKeown 87]) that instruct the planner where in the representation network to search for additional sentence topics, given an initial topic and such goals as to describe it, to compare and contrast it, etc. In similar fashion (as described in chapters 2 and 4), PAULINE has three plans -- the DESCRIBE plan to find descriptive aspects of objects, the RELATE plan to relate events and state-changes, and the CONVINCe plan to help slant the text in order to convince the hearer of some opinion. Whenever it performs topic

collection, PAULINE applies the prescriptive steps of the appropriate collection plan to each candidate topic, and then in turn to the newly-found candidate topics, for as long as its pragmatic criteria (amongst others, the values of **RG:haste** and **RG:openmindedness**) allow.

Clearly, the details of these plans depend very much on the representation scheme used. For more detail, consider the **good results** strategy, one of the steps of the **CONVINCE** plan:

```

COLLECT, from the topic,
    - all its direct RESULTS
    - if it is a MOP, its 'INAL RESULTS
    - if it is part of a plan or otherwise serves a goal, the
      DESIRED STATE of the goal
FOR each of these results,
    IF it carries the affect GOOD for the speaker
    THEN retain it as a candidate topic
  
```

For example, PAULINE uses the **CONVINCE** plan in generating the shantytown examples as follows: When speaking as a protester, the instances of support by the local community and faculty members are obviously **GOOD** and are included, whereas the fact that Yale was at first lenient and allowed the shantytown to remain in place until the meeting of the Yale Corporation is **BAD** and is ignored. Also, the arrest of 76 students is **GOOD** for this point of view and is included, but the commission of investigation is not. The contrary is true when PAULINE is a Yale official.

7.4.5 Topic Interpretation

As described in chapter 3, PAULINE has the ability to perform limited inference in order to interpret its input topics in ways that help achieve its pragmatic goals. These inference strategies are defined as prescriptive plans. Activated rhetorical goals of opinion (such as the

goal to interpret the opponent's actions as confrontational) cause interpretation inferences to be applied to the collected candidate topics during topic organization.

In order to interpret the input topics as instances of some concept, the interpretation process must recognize when the topics (or some of them) conform to the definition (or part of the definition) of the concept. Thus, either concepts must be defined in such a way as to allow a general process to read their definitions, or inferences must exist that fire when a definition is matched -- in other words, whose left-hand sides are the definitions and whose right-hand sides assert the existence of the concept. PAULINE is implemented with the second approach. Bottom-up interpretation inferences reside in memory and the lexicon as part of the definitions of concept types; top-down interpretation inferences are contained in plans that serve the pragmatic goals.

Interpretation inferences are defined as configurations of concept types and relations among them. A configuration is the description of the way in which a collection of concepts must relate to each other to form a legitimate instance of another concept. The interpretation mechanism matches candidate configurations against its collected topics, and, if matched, creates a new instance of the interpretation and adds it to memory. The program can then generate text from the interpretation instead.

Each configuration is a pattern, in the form of a list of triplets (*type ?var pattern*), where:

- *type* is either the type (in the property inheritance memory network) of the concept currently to be matched, or a variable *?var* which must have been encountered before
- *?var* is either (), or a variable *?var* by which the current concept will be identified later in the match, or a number of such variables that must be bound to different concepts for a match
- *pattern* is a list of (*aspect config*) pairs, where the filler of each *aspect* must recursively match the *config*, which is again a pattern

Configuration patterns obviously depend on the exact representations used. For example, the configuration for the concept BEAT is:

```
(VOTE-OUTCOME ?X                                ; ?X is a VOTE-OUTCOME
  (instance (ELECTION ?Y))                        ; in some primary ?Y,
  (relations (REL-GREATER ()                     ; and it is greater
              (conc1 (?X))                       ; than
              (conc2 (VOTE-OUTCOME ()            ; another VOTE-OUTCOME
                      (instance (?Y))))))        ; in primary ?Y
```

which means: some concept is of type VOTE-OUTCOME; its aspect RELATIONS contains a GREATER relation, of which the greater part is that same concept and the smaller part is another VOTE-OUTCOME in the same primary. Thus, since Kennedy's outcome is the outcome of a primary and it is greater than Carter's outcome, the two form an instance of BEATING. Most configurations are considerably more complex; consider, for example, CONCILIATION:

```
; ?X's action serves one of his own goals
(list T ()                                     ; The input is any concept (say,
  (list 'actor                                ; the permission to reconstruct)
    (list AGENT '?X))                        ; and its actor is ?X (Yale).
  (list 'relations                            ; This concept
    (list REL-SUBGOAL-TO ()                  ; serves a goal
      (list 'conc2
        (list HAVE-GOAL ()
          (list 'actor                        ; held by ?X (the goal that
            (list '?X))))))                 ; the students be orderly).
    ; and also serves ?P's (someone else's) goal
  (list 'relations                            ; Also,
    (list REL-SUBGOAL-TO ()                  ; this concept serves a goal
      (list 'conc2
```

```

(list HAVE-GOAL ()
  (list 'actor
    (list AGENT '?P          ; held by ?P (the students)
      (list 'opposites:      ; who is not the same
        (list '?P '?X))))   ; person as ?X.
    ; which opposes another goal that ?X has
    (list 'relations
      (list REL-OPPOSING () ; Furthermore, this goal opposes
        (list 'conc2
          (list HAVE-GOAL () ; another goal held by ?X
            (list 'actor      ; (the removal of the shanties).
              (list AGENT '?X)))))))))))))

```

7.4.6 Additional Topic Inclusion

During the course of topic organization, the generator may find additional candidate topics. Whether or not to include these instances can only be decided when such topics are found; the relevant strategies are therefore restrictive. For example, after interpreting input topics as an instance of some new concept, the new concept may (in a memory organized after the style described in [Schank 82]) furnish other instances, which could be included using phrases such as "that reminds me...".

Under certain circumstances, explicit statements of opinion can be included in the text. Rather than give the generator the prescriptive goal to create such circumstances, which may be very difficult, restrictive plans enable their inclusion to be considered whenever the circumstances arise. These restrictive plans also serve rhetorical goals of opinion, as described above and in chapter 2. For example, under appropriate circumstances, PAULINE has the (restrictive) option to include explicit statements of opinion before sentences, such as "it pisses me off", "I am happy", and "X wanted to be reasonable".

7.4.7 Topic Juxtaposition

As described in chapter 4, whenever two topics are being generated, the generator can search for suitable multi-predicate phrases (such as "not only X, but Y") in which to frame them in order to achieve desired opinion-related and stylistic effects. The task of topic juxtaposition is best implemented restrictively by presenting the candidate topics as options to strategies that check the restrictions on the use of phrases and select suitable ones. (The equivalent prescriptive implementation amounts to giving the program goals such as [FIND IN MEMORY TWO TOPICS THAT FIT INTO A *not only but* PHRASE], a much less tractable task.)

When PAULINE has the goal to convince the hearer of its affect, and when the rhetorical goals **RG:formality** and **RG:simplicity** do not prohibit long and complex sentences, the program plans phrasal juxtaposition in the following way: First, it determines the hearer's affect for each of the current candidate concepts. By the affect rule, GOOD topics can be said; NEUTRAL ones should be enhanced to look GOOD; and BAD ones should be avoided or mitigated:

1. determine the hearer's affect for the current candidate topics
2. cast BAD topics in mitigator phrases with GOOD topics; cast NEUTRAL topics in enhancer phrases with GOOD topics
3. say additional GOOD and NEUTRAL topics; suppress unmatched BAD ones

When there are too few GOOD topics to match all the BAD ones, some of the latter must be mitigated by other means or not said at all. This decision depends on the rhetorical goal **RG:timidity**: a *reckless* approach permits unmatched BAD topics to be said, hoping that other generator decisions can mitigate them, as described below. A *timid* approach suppresses unmatched BAD topics.

7.4.8 Topic Ordering

The ordering of topics in the paragraph is best achieved prescriptively. Different circumstances call for different orderings; newspaper articles, for instance, often contain an introductory summarizing sentence, as in example 1, chapter 1. In contrast to the schemata used by McKeown, steps in PAULINE's topic collection plans are not ordered; additional plans must be run to ensure coherent text flow. PAULINE uses one of two topic-ordering plans that determine the position in the paragraph of an introductory sentence, descriptive and supporting sentences, etc.

7.4.9 Syntax

The goal to produce grammatical text underlies the design of all generators (except perhaps [Clippinger 74]), and causes the (prescriptive) use of grammatical rules and phrases. The syntax-level realization decisions are described in chapter 6. Since, however, syntax and semantics are underspecific — one can express the same semantic information in many ways — pragmatic considerations provide additional criteria for making selections. As argued in chapter 2, pragmatic information is communicated through appropriate slant and style.

(a) **Sentence Slant:** In addition to the slanting techniques mentioned above (appropriate interpretation, opinion inclusion, and topic juxtaposition), a number of other techniques exist to help slant text. As described in chapter 4, these techniques include the use of stress words, adjectives, adverbs, verbs that require idiosyncratic predicate contents, nouns, etc. As with style, due to the local (i.e., non-formative) nature of most of these techniques, they are best implemented restrictively.

As described in chapter 4, a number of rules exist for choosing sentence subjects in order to produce flowing, natural text. PAULINE does not use the notion of focus, but it does use affect as a criterion for subject choice; its two rules are (the second is the equivalent of McKeown's default rule, [McKeown 82]):

1. select the subject from candidates with GOOD affect for sentences with GOOD affect and from candidates with BAD affect for sentences with BAD affect
2. otherwise, select as subject the sentence topic's AGENT, if it exists

For example, when PAULINE is defending Jim, it says:

(a) MIKE ONLY WAS BUMPED ONCE BY JIM

In order to build appropriate sentences, the generator must inspect the formative information it uses — whether phrases, idiosyncratic predicate forms for verbs, or noun groups — for the default implications of its parts. Just as the mitigator phrase “X; however, Y” imputes to Y the opposite affect of X, the predicate forms of verbs can be viewed as having fields with characteristic affects. For example, in the form for “beat”,

“A beat B [in C]”

before anything is known about the field fillers, A carries the default affect GOOD, B BAD, and C GOOD. These field affects can be used to impute affect, since the relative polarity of the fields remains unchanged: B must carry opposite affect to A and C. Seen from the perspective of verb choice, if the speaker likes Pete, he should not use “beat” (unless Bill and/or the race have previously been established as BAD):

“[That blighter] Bill beat [my hero] Pete in the [crooked] race”

Similarly, he should not use “lose” either, since in

“A lost [to B] [in C]”

field A carries BAD to fields B's and C's GOOD affects. However, he could use “win”, which doesn't require a direct object, or “get”, which avoids the losing altogether:

“Bill won the race”

“Carter got 20515 delegates”

and enables him to avoid placing his sympathy in a BAD field. This is what PAULINE does to produce (a) rather than (b) when it supports Carter in the example above. The incisive effect of this strategy (still supporting Carter) is even clearer when the sentence is embedded in a phrase:

(b) **ALTHOUGH KENNEDY WON THE PRIMARY, CARTER IS AHEAD**

(c) **ALTHOUGH KENNEDY BEAT CARTER IN THE PRIMARY, CARTER IS AHEAD**

(b) **Sentence Style:** In addition to being slanted, text can be hasty, formal, forceful, etc., as described in chapter 5. Control of style is pre-eminently a restrictive task, since the relevant selections typically occur at relatively advanced points in the generation process and hence have relatively local effect: points such as deciding on conjunction and relativization (i.e., length of sentences); determining the presence and position of adverbial clauses, adverbs, and adjectives; selecting specific words (especially verbs and nouns). When the realizer reaches a decision point between options that are equivalent with respect to semantics and opinion (slant), restrictive plans that achieve stylistic effects come into play.

7.4.10 The Planning Process, Step by Step

In summary, then, PAULINE's topic organization (planning) phase proceeds in the following way:

1. the user gives PAULINE one or more topics, which are part in the representation of the episode, embedded in the concept representation network. For example, in the shantytown examples, PAULINE is given three representation elements: the construction of the shantytown, its removal, and the subsequent permission to have it rebuilt
2. next, the user inputs the characteristics of the speaker, the hearer, and the conversational setting. This includes the interlocutors' sympathies and antipathies, and values chosen from the options listed under PAULINE's characterization of the pragmatic aspects of the conversation, chapter 2.3
3. finally, PAULINE is given interpersonal goals with respect to the hearer's opinions, knowledge, relationship to the speaker, etc. These values are also chosen from the options listed in chapter 2.3

4. **PAULINE** activates the rhetorical goals with appropriate values, using the activation rules given in chapters 3.6, 3.7.2, 4.6.1, 4.6.2, 5.2, 5.3, and 5.4
5. if permitted by **RG:haste** and **RG:openmindedness**, a topic collection plan (such as the **CONVINCE** plan) is activated and its steps applied to the next topic
6. when enough collection has occurred, the program starts topic organization
7. when **RG:simplicity** and **RG:haste** permit, a topic ordering plan is run and the topics appropriate for the next stage of the paragraph are gathered as the current candidate topics, as described in chapter 4.4.2
8. if **RG:detail** and **RG:haste** permit, top-down interpretation occurs: from the interpretation inferences associated with the activated rhetorical goals of opinion, the planner collects all the configurations that match candidate topics, as described in chapter 3 and above
9. also, if those goals permit, bottom-up interpretation occurs: using each candidate (and sometimes related ones) as a pivot, the planner searches for back-pointers to configurations and collects those that match the candidates
10. a suitable configuration is selected, as described in chapter 3.4
11. if not yet in memory, the planner builds a new interpretation of the matched candidates and indexes in memory it off the interpretation concept
12. if the newly-found interpretations contains specific instances, these are the reminders; a suitable reminding is selected, subject to **RG:color** (chapter 3.4)
13. the reminding probably matches the candidates only in some aspects — the aspects that match the configuration — and probably has other aspects — say, information about events whose equivalents have not yet occurred — which may be of interest to further generation. The rhetorical strategies construct, from the input, hypothetical equivalents of the reminding and examine them to see whether to say the reminding

14. the matched candidates and the reminding are packaged into a syntax goal (the goal to say them)
15. the interpretation process is repeated, subject to **RG:haste** and **RG:detail**
16. next, topic organization proceeds, again subject to **RG:haste** and **RG:simplicity**
17. depending on **RG:timidity** and **RG:partiality**, the planner tries to embed one or more of its topics (each in a syntax goal) into multi-predicate enhancer or mitigator phrases
18. depending on **RG:simplicity**, **RG:formality**, and **RG:haste**, the planner tries to embed one or more of its topics (each in a syntax goal) into phrases that express goal, spatial, temporal, etc., relations
19. this topic organization process is repeated, depending on **RG:haste**
20. the rhetorical planner finds an appropriate syntax specialist for each goal's topic and adds it to the syntax goal
21. the goals are dispatched to be generated

An annotated trace of a PAULINE session appears in appendix A.

7.5 Conclusion

The selections distributed throughout the generation process are not just a set of unrelated ad hoc decisions; they are related in ways that permit the creation of style and slant. Therefore, they require control. Since, however, traditional top-down prescriptive planning is unable to provide this control, a different *kind* of planning that is bottom-up is required. With this approach, we can identify the tasks that require such planning as restrictive and build plans and strategies to control them in a uniform fashion. This enables us to interleave realization with the most appropriate style of planning, either prescriptive or

restrictive, and the result, limited-commitment planning, greatly enhances our ability to build generators that produce good, pragmatically appropriate text.

Restrictive planning raises a number of interesting issues dealing with conflict resolution among goals that constantly compete for resources. Two such issues are goal status monitoring and the determination and variation of relative goal priority. These issues require further investigation, in both the generation and the planning research.

Chapter 8

A Review of Generation Literature

8.1 The Three Questions of Generation

The study of natural language generation by computer has traditionally been divided into two questions: *what shall I say?* and *how shall I say it?*. Philosophers of language have done much work on the former question; the latter has been attacked mostly by linguists. Since linguists tend to produce more easily encodable results than philosophers, early generation work in AI concentrated on the latter question, with the result that today, using available knowledge, one can quite easily build a useful (though limited) sentence generator program, but one cannot so easily build a program that selects topics and performs the planning required to organize them into coherent text.

The early programs, as well as many existing ones, do not address the *what* question -- the issues of topic selection and organization. For these generators the task is simply: someone (or some program) feeds them a chunk of representation, and they find a way of saying it -- the whole chunk and nothing else. Their answer to the *what* question is simply *say it all!*. On the other hand, the few programs that do attempt topic selection and

organization tend to skim when it comes to the *how* question; in fact, some do not even pretend to make sentences, and most of the others use "canned" language — templates. In addition, a number of generator tasks, such as interpretation and multi-predicate sentence construction, have never really been attempted by any program, because they lie between the two questions, and are not properly addressed by either of them alone.

Putting the two sets of tasks together is difficult for a number of reasons — the interaction between them is not understood, and neither is their relative timing. In addition, combining them intensifies the sheer difficulty of guiding the generator through the large number of topic-related and syntactic opportunities that exist in language, because when implemented separately, each task usually assumes that the other does a lot of the hard work. That is, topic collection programs typically expect that their (assumed) syntactic components will be able to make good, coherent, flowing text with whatever is handed to them, and language producer programs assume that their (assumed) topic collection components will perform a large amount of difficult but necessary structuring on the input topics.

When you address this problem of guidance a third question arises: *why should I say it?* The thesis of this dissertation is that the answer to this question is, to a large extent, pragmatics. The generator's goals, both informative and pragmatic, help to differentiate among its options and help indicate which selections should be made and which should be avoided. But since the *why* question is very complex, more difficult to answer than either the *what* or the *how* questions, it has largely been avoided to date. But we cannot avoid it for ever. If we want programs to produce natural, high-quality text, we have to take pragmatics into account.

The first part of this chapter contains a brief overview of AI work in language generation, and indicates how PAULINE fits in. The second part discusses what lessons can be learned from this attempt at incorporating pragmatics into generation, and what should be done next.

8.1.1 How shall I say it?

Transformational Grammar:

When given the task to build a generator, transformational grammars seem at first glance to be eminently suitable for answering the *how* question. However, they are not: in fact, TG was originally meant to be purely descriptive, not functional:

...it seems absurd to suppose that the speaker first forms a generalized Phrase-marker by base rules and then tests it for well-formedness by applying transformational rules to see if it gives, finally, a well-formed sentence. But this absurdity is simply a corollary to the deeper absurdity of regarding the system of generative rules as a point-by-point model for the actual construction of a sentence by a speaker.

[Chomsky 65, p 139]

If neither the psychological validity of the generator nor its efficiency are important issues, then transformational grammar can be a basis for building a generator.

Template Generators:

For small, practical domains, probably the easiest and most-used generators are template generators. A template is associated with each representation element type, in such a way that the templates can nest in the way representation elements nest. Generation then becomes a matter of filling in the blanks. Template-based generators can produce very impressive text, but, of course, are hardly flexible. For a typical example, see [Swartout 81]. More sophisticated template generation is performed by [Rösner 87] and [McDonald & Woolf 85].

Augmented Transition Networks:

In generation work of a fifteen years ago no text planning phase ever appeared. Typically, an ATN was defined to embody some subpart of the grammar of the target language. It was given representation elements defined in a semantic-type network. Under guidance of

the input, the ATN network was traversed until some path was found that ended in an end state of the ATN and that exhausted the input; the traversal route then provided the sentence. In this way the linearization of the input can be achieved in simple and highly unhumanlike ways. ATNs, for example, cannot make mistakes. But the technology is clean, easy to write and extend, and works. For an excellent description of an ATN generator, see [Simmons & Slocum 72].

Another example of an ATN generator is Goldman's BABEL ([Goldman 75]). One of the nice aspects of BABEL was that it could express paraphrases of a sentence to varying levels of precision: "John strangled Mary" could be said as "John put his hands around Mary's neck and choked her" or as "John put his hands around Mary's neck and squeezed her windpipe shut and caused her not to be able to breathe which caused her to die". However, BABEL had no idea when each paraphrase was appropriate! The answer, clearly, is based upon a *why* question.

ATNs are well suited for storing grammatical information; they are not, however, suited for storing all the other information a generator requires. For example, when more than one word exists for a piece of the input, other processing is needed. BABEL used discrimination nets with semantic tests in order to find words for its input elements. For example, when confronted with the Conceptual Dependency element INGEST (see [Schank 75, 78]), the program inspected the nature of the OBJECT role filler; if it was a liquid, the word chosen was "drink", if a solid, "eat", and if a gas, "breathe".

Other Ways of Constructing Syntactic Representations:

A number of more recent generators use rules to construct explicit syntactic representations of sentences; in a second pass, they substitute words and produce the text. For example, Tree Adjoining Grammars (see [Joshi 87], [Kroch & Joshi 86]) contain rules for constructing and embedding syntax trees in ways that produce valid sentences. Wong describes a more limited generator in [Wong 75].

The most advanced generator of this type is MUMBLE ([McDonald 80, 81], [McDonald & Pustejovsky 85]). Entities called *realization classes* (corresponding roughly to

PAULINE's syntax specialists) communicate using an elaborate protocol to construct complex syntax trees that are then uttered after a second pass lexical substitution phase. McDonald takes seriously the position of autonomous syntax: this position has the advantage that the boundaries of the generator are well-defined and that, in time, a program can be constructed that embodies most of (or all!) the syntax of English. But this position also has a disadvantage, one that outweighs the advantage: being autonomous means *not relying* on outside influences, meaning that once the generator starts working on its input it runs alone. Since clearly pragmatic and semantic issues influence the generation of text, this position requires that these issues' effects be somehow already incorporated in the input to the generator. Therefore the input must be prestructured to an extent that *really* makes the generator's decisions uninteresting: if the decisions are not allowed to carry semantic or pragmatic import, what then *are* they allowed to do? Pragmatic issues *cannot* be divorced from low-level generation; this dissertation has hopefully demonstrated this fact, if nothing else.

Work on the Lexicon:

A number of people have investigated ways of representing syntactic information in the lexicon. In [Danlos 85, 87], Danlos develops a phrasal grammar (similar to but richer than that described in chapter 6 and appendix B) by comparing linguistic issues in French and English. She justifies each syntax specialist and identifies syntactic and semantic conditions under which specialists are appropriate. In [Jacobs 85, 87], Jacobs develops a representation language for a phrasal lexicon that captures syntactic and semantic commonalities in language, enabling quasi-metaphorical expressions such as "Frazier gave Ali a punch" and "Ali took a punch from Frazier". [De Smedt & Kempen 87] are investigating incremental generation with a program in which conceptual input can be added to the partially realized tree of linguistic knowledge. This work draws on the psycholinguistic work of [Kempen 76 and Kempen & Hoenkamp 78].

Systemic Grammar:

Based on the ideas of Halliday (see, e.g., [Halliday 76, 78] and [Fawcett 80]), a number of systemic generators have been built. As with ATNS, systemic grammars contain the structure of the grammar embodied in a network. A systemic grammar is a large collection of dis-

crimination nets called systems, where each system (one or more discriminations) computes the value(s) of one or more grammatical aspects such as number and mood. Collections of systems thus perform the functions of the syntax specialists. However, unlike phrase structure grammars (which include all the abovementioned implementations), systemics or collections of systems do not have to be activated sequentially — they can become active and perform their tasks as soon as their input conditions are satisfied. This fact makes systemic implementations potentially far more powerful than phrase structure grammars.

Insofar as syntax (and pragmatics; see [Fawcett 80]) can be represented as collections of aspects whose values can be computed via discriminations, systemic grammar offers a clean and powerful formalism with which to build a generator. Whether all the information required for generation can in fact be represented in discrimination networks is an open question. [Patten & Ritchie 87] give a formal definition of systemic grammar; the generator in [Davey 79] describes tic-tac-toe games. Penman, the world's largest automated systemic generator, is described in [Mann 82, 83a, 83b].

8.1.2 What shall I say?

The second question of language generation, *what shall I say?*, has received less attention. Intrinsically, it is a much harder question to answer, because it depends on factors about which the speaker can never have complete knowledge, such as the hearer's knowledge and beliefs. It is much easier always to make grammatically correct sentences than always to say appropriate things!

In the simplest text planning systems, the planners make only very high-level decisions and play no further role in the realization process. One example is Cohen's program (Cohen 78) that reasons about the user's knowledge state in order to decide on an appropriate speech act. The program included a realization component.

Scriptal Planning

A major planning task in the construction of a paragraph is the collection and ordering of

sentence topics. In [McKeown 82], McKeown made the (with hindsight, rather obvious) breakthrough: she defined four so-called schemas which are, in essence, generation scripts that direct the topic collection process and (together with rules of topic flow) control the ordering of sentence topics. Thus, for example, the IDENTIFICATION schema prescribes first saying the object's name and type, then giving (some of) its attributes, then its uses, specializations or instances, etc. Her other schemata are ATTRIBUTIVE, CONSTITUENTS, and COMPARE & CONTRAST. Her generator produced output for a database Q/A system that contained information about war machines.

Being script-like, these schemas are tailored to specific circumstances. For hearers with different levels of knowledge, schemas must prescribe appropriate details of the topic. This problem is being investigated by Paris (see [Paris 85] and [Paris & McKeown 87]). For hearers with different interests, different aspects of the central topic are relevant. McCoy ([McCoy 85, 87]) uses different schemas to correct different hearer misconceptions in different ways. She describes using a *salience value* for aspects of topics and in schemas to help decide what additional sentence topics the schemas should collect. In her system, the salience values reflect characteristics of the hearer and are determined by the hearer's questions.

A variety of other special-purpose generators that use strategies in a script-like way exist, for example, [Kukich 83] describes the generation of stock market reports, [Novak 85, 87] develops strategies for describing relative motions of objects, [Reithinger 84, 87] generates referring expressions and answers in interactive systems. In [Woolf 84] and [Woolf & McDonald 84], Woolf describes a program that tutors students, by planning which schema to select and when to change it, based on its representation of their knowledge and misconceptions.

Hierarchical Expansion Planning:

In systems with more flexible plans, additional work must be done before a text realization can be started. Mann's Rhetorical Structure Theory ([Mann 84], [Mann & Thompson 83, 87]) is an attempt to build up a set of rhetorical structures that can be used to construct paragraph-length and longer arguments. This theory is still being constructed. Appelt's

planner KAMP ([Appelt 81, 82, 85]) is the most general and powerful text planning program yet developed, being an implementation of part of the general-purpose planner NOAH ([Sacerdoti 77]). KAMP starts with a number of goals to inform the hearer of a number of facts, reasons about the hearer's knowledge state, and constructs utterances that contain an appropriate amount of information in an appropriate order. In [Appelt 81] the impression is given that the program introduces grammatical constraints midway in the plan expansion process and finally works exclusively with syntactic constructs to produce the text. This turns out to be difficult to do, on p. 113, Appelt expresses dissatisfaction with this way of integrating grammatical and other planning constraints on the grounds that (a) the grammar cannot easily be evaluated without running the program and (b) the grammar cannot easily be altered. In practise, however, KAMP initially used simple templates for the actual text forms and later used a functional unification grammar back end ([Appelt, personal communication, 1986]). In subsequent work, Appelt has split off the syntactic component and uses KAMP simply for the "normal" planning activities.

General Problems with Planning:

In almost all existing text planners, information flows one-way, from the planner to the realizer ([McDonald, personal communication, 1984], [McDonald 80], [McDonald & Pustejovsky 85], [Appelt 81]). In these systems, two approaches are possible: either all requisite information is pre-planned before realization starts, or planning and realization are interleaved and the planner continues whenever the realization phase completes a subtask. Appelt followed the former option; McDonald the latter. Pre-planning the requisite information is comparatively easy if the criteria for these decisions are based on relatively uniform grounds (for example, syntactic grounds only, including notions such as sentence focus and stress), since they are then relatively simple and only impinge on expansion at a small number of points, such as subject, predicate, tense, content only. This is one reason why Appelt's program KAMP could operate as it did: it took into account only the hearer's knowledge state and never had the support of free writing or varying variations of text.

When less homogeneous and more complex information is taken into account, however, pre-planning all the likely decisions and generating many hypotheses is an overwhelming task — the

planner must effectively span the space of possible locutions — which is not a practicable solution to the problem. In later work, McDonald argues together with the psycholinguists Levelt ([Levelt & Schriefers 87]) and Bock ([Bock 82, 87]) that the two stages, planning and realization, interface at the so-called *message* level. This facilitates limited-commitment planning and two-way communication between the modules, as argued in [Hovy 85] and in chapter 7: typically, requests for guidance together with the available options pass from the realization component to the planner, and the planner's selections and decisions are handed back to the realizer. This is the way PAULINE is implemented.

Other Planning-Related Work:

Much work has been done on the planning and generation specifically of multi-sentential units of text such as conversations. [Grosz 77, 80] and [Sidner 79] discuss rules of focus shift, topic flow, and anaphora; in later work they collaborate and are currently developing a theory of conversations using notions such as speaker intention and attention ([Grosz 85], [Grosz & Sidner 86]). [Reichman 81] constructs a grammar for the structure of discourse.

8.1.3 Why should I say it?

To date, only a handful of programs construct, in answer to this question, explicitly represented speaker goals to motivate the decisions they make. Jameson [Jameson 87] is building a system that constructs its text by selecting utterances from scripts of templates, based on the desired pragmatic import of the final text. Clippinger ([Clippinger 74, 75]) built a complicated system ERMA that modelled one paragraph of text produced by a person in a psychiatric session. ERMA consisted of five modules: CALVIN (topic collection and organizing), MACHIAVELLI (topic organization, phrasing, etc.), CICERO (realization), FREUD (monitoring the origins of generator tasks), and LIBNITZ (concept definition network). Different parts of the program had characteristic effects on the text, and the hope was that ERMA's text would produce some of the false starts, hesitations, suppressions, etc., that people make when they speak. Since ERMA was *built* to model the paragraph of text, it is not, of course, however, though the program is interesting and four of the five parts believable (perhaps, with more examples, the function of FREUD would be clearer), the overall

theoretical contribution of ERMA is not clear. Certainly we are a long way from being able to assess all the claims made about the work.

This brings us to PAULINE (which has been called "a parameterization of ERMA" [McDonald, personal communication, 86]). PAULINE was inspired by the work of McGuire (see [McGuire 75]), who first worked on what eventually became the Carter-Kennedy example. PAULINE was built to illustrate the utility of pragmatic information in making planning and realization decisions that are otherwise not decidable.

8.2 Conclusion: Doing it Better Next Time

One of the strong points of PAULINE is the use of a phrasal lexicon, in which, as described in chapter 6, words, phrases, and formative syntactic information associated with phrase structure symbols are contained in homogeneous form. The lexicon should, however, contain even more information. It is a serious handicap to the planner that lexicon elements do not, for the most part, contain explicitly represented rhetorical information. That is to say, for example, that the phrasal pattern "not only X but Y" does not explicitly provide the planner the facts that it is *moderately formal*, *quite complex* (containing two predicates, and inverting word order), and an *affect-enhancer* in which the first filler, X, must hold some affect for the speaker. Currently, much of this type of information is procedurally encoded into discrimination nets, something that makes both planning and the addition of new lexicon elements difficult.

Adding all such information to the lexicon explicitly requires, first, the formulation of a set of terms, based on the program's rhetorical goals and strategies, which the planner can manipulate to construct plans, and second, the reformulation of the planner and monitor (described in chapter 7). As it is, the planner is extremely straightforward: it simply tries to achieve each activated rhetorical goal, preferring, in conflicts, the goals least satisfied so far. It has no record of the number of times a rhetorical goal was actually thwarted or blocked, and is thus unable to recognize the need to repair mistakes. Including rhetorical information explicitly in the lexicon — both about goals achieved and about goals thwarted

— will enable the planner to know when to perform repairs. Of course, the planner must then also have repair strategies and be able to decide what strategies are required. This problem is now being attacked in the general planning and execution monitoring literature in AI; see, for example, [Doyle, Atkinson & Doshi 86]). The issue of limited-commitment planning under multiple goals is a very interesting research topic.

The most obvious shortcoming of the theory presented here is the lack of sophistication of the pragmatic features used in PAULINE. Each feature — hearer knowledge, opinion, emotional interpersonal relation, relative social status — is a field of study in itself. However, this is not the most serious shortcoming, because the rhetorical goals posited here provide a satisfactory level of organization of information, midway between the airy world of pragmatics and the concrete requirements of generators. That is to say, any generator that can realize a given input in different ways requires reasons for choosing one way over another, and these reasons will always pertain to the stylistic and opinion-based considerations — the rhetorical goals and strategies — discussed in this work. On the other hand, any system that manipulates information of pragmatic nature and wishes to use this information to guide a generator will always do so using intermediate strategies that correspond to the rhetorical goals. The clumsiness of the rhetorical strategies and of the goals' activation conditions, as described here, is due to the unsophisticated nature of the terms used to characterize pragmatic information and generator decisions; increasing the level of sophistication, and experimenting with combinations of strategies to form different rhetorical goals, should be the focus of continued research.

Appendix A

A Short Annotated Example

In order to illustrate the interleaving of prescriptive and restrictive planning and text realization, this appendix contains an annotated trace of PAULINE's generation of a shantytown text. In the example, PAULINE is a protester trying to convince its audience at a meeting; though it is somewhat formal, it is not pressed for time:

AS A REMINDER TO YALE UNIVERSITY TO DIVEST FROM COMPANIES
DOING BUSINESS IN SOUTH AFRICA, A LARGE NUMBER OF STUDENTS
CONSTRUCTED A SHANTYTOWN --- NAMED WINNIE MANDELA CITY --- ON
BEINECKE PLAZA IN EARLY APRIL. AT 5:30 AM ON APRIL 14, YALE HAD
OFFICIALS DESTROY IT; ALSO, AT THAT TIME, THE POLICE ARRESTED 76
STUDENTS. A LARGE NUMBER OF LOCAL POLITICIANS AND FACULTY MEMBERS
EXPRESSED CRITICISM OF YALE'S ACTION. FINALLY, THE UNIVERSITY
GAVE THE STUDENTS PERMISSION TO REASSEMBLE IT.

Initially the conversation conditions are set up. The input topics are the building of the shantytown (CONSTRUCT-1), the demolishing (CONSTRUCT-2), and Yale's permission for its reconstruction (MTRANS-8). As sympathies, PAULINE is given the protesters' goal to have Yale divest from companies with business in South Africa (GOAL-1). The hearer, a predefined pro-Yale person, has as sympathies Yale's goals to keep the university orderly (GOAL-2) and to be reasonable in the face of argument (GOAL-3):

SPEAKER'S INTERPERSONAL GOALS:

Speaker's conversation with UNIV

about (CONSTRUCT-1 CONSTRUCT-2 MTRANS-8)

CONVERSATION SETTING:

Setting:

- Available time: MUCH
- Tone: FORMAL
- Conditions (noise, etc.): GOOD

Speaker:

- Sympathies: (GOAL-1)
- Antipathies: (GOAL-2 GOAL-3)
- Knowledge level: EXPERT
- Interest level: STRONG
- Emotional state: NORMAL

Hearer:

- Sympathies: (GOAL-2 GOAL-3)
- Antipathies: (GOAL-1)
- Knowledge level: STUDENT
- Interest level: SOME
- Emotional state: CALM
- Language ability: NORMAL

Relationship:

- Extent of acquaintance: STRANGER
- Hearer's social status relative to Speaker: EQUAL
- Emotion between interlocutors: LIKE

SPEAKER'S GOALS:

- Topic collection plan: *CONVINCE-PLAN*
- Conversation time: MUCH

Desired effect on Hearer's Knowledge and Beliefs:

- Hearer knowledge level: TEACH-TALK
- Hearer opinions: SWITCH

Desired effect on Hearer's Goals and Emotions:

- Alter Hearer's goals: SWITCH
- Hearer's emotional state: NORMAL

Desired effect on Hearer's Relationship to Speaker:

- Distance between interlocutors: NORMAL
- Hearer emotion toward Speaker: LIKE
- Relative status between interlocutors: EQUAL

Next, of course, the rhetorical goals are activated:

Setting up rhetorical goals from interpersonal goals...

SPEAKER: PROTESTER

HEARER: YALIE

SPEAKER'S RHETORICAL GOALS:

AFFECT GOAL -- *CONVINCE-GOAL*	KNOWLEDGE GOAL -- ()
TOPIC SLANT PLAN -- *OPINION-PLAN*	TOPIC COLLECTION PLAN -- *CONVINCE-PLAN*
FORMALITY -- HI	DETAIL -- MED
FORCE -- MED	HASTE -- MED-LO
TIMIDITY -- LO	PARTIALITY -- MED
SPEAKER-REF -- HI	HEARER-REF -- MED
COLOR -- HI	FLORIDITY -- MED
RESPECT -- MED-HI	SIMPLICITY -- MED
WARMTH -- HI	VERBOSITY -- MED
AGGRESSION -- MED	INCITEMENT -- MED

In order to elaborate on the input and produce better text, the program can search for additional topics. Of course, such topics will be related to the input in some way; thus to know where to find likely candidates, PAULINE follows the prescriptive instructions of one of its topic-collection plans. As described in chapter 4, since its sympathies differ from those of the hearer, the program selects the CONVINCE plan. (First, however, it must check whether (a) it has time to do topic collection, and (b) whether the hearer's opinions about some aspect of the topic disagree with its own, otherwise there is nothing to dispute!):

Topic collector considering new topic: CONSTRUCT-1

PRAGS: level of HASTE for extent of topic collection

-->checking rhet goal: RG:HASTE

Searching for a topic-expansion plan for goal CONVINCE-GOAL

Deciding whether to convince about CONSTRUCT-1's affect

-- checking whether I agree with hearer on all its aspects

Working out *SELF*'s affect for SHANTYTOWN-1 ...none
 Working out *UNIV*'s affect for SHANTYTOWN-1 ...none
 Working out *SELF*'s affect for STUDENTS-1 ...none
 Working out *UNIV*'s affect for STUDENTS-1 ...none
 Working out *SELF*'s affect for SBG:CONS1-G1 ...GOOD
 Working out *UNIV*'s affect for SBG:CONS1-G1 ...BAD
 ... no: we differ about its RELATIONS (SBG:CONS1-G1)
 Checking amount to be said ... there is enough to say
 CONVINCE-PLAN is appropriate

Now applying *CONVINCE-PLAN*'s steps to CONSTRUCT-1

- trying plan step to find worse example
- trying plan step to find good (or not bad) result
- trying plan step to find good goal
 - found the following: -> GOAL-1 is the SUBGOAL-TO of CONSTRUCT-1
- trying plan step to find good (or not bad) relations to concepts
 - found the following:
 - > GOAL-1 is a SUBGOAL-TO relation of CONSTRUCT-1
 - > SUPPORT-1 is a RESULT relation of CONSTRUCT-1
- trying plan step to find good side-effect
- trying plan step to find appeal to authority
- trying plan step to find suitably slanted affect
 - found the following: -> CONSTRUCT-1 can be said with suitable slant

CONSTRUCT-1 spawned the following topics: (SUPPORT-1 GOAL-1 CONSTRUCT-1)

Since the speaker and hearer have different opinions about GOAL-1 (the protesters' goal to have Yale divest), the CONVINCE plan can be applied to the input topic; this directed search produces two additional topics: GOAL-1 and SUPPORT-1 (the support given the protesters by the local community).

At this point, CONSTRUCT-1 has been examined and can be said directly -- that is, the realization phase can begin. Indeed, if the rhetorical goal **RG:haste** had the value *pressured*, this is what would happen, although this would mean that the program could not perform further topic organization tasks (such as interpretation and phrasal juxtaposition) since it would have no additional topics. This is the first point at which interleaving between planning and realization could occur. However, since the program is not hasty,

it can continue the search for additional topics in order to facilitate such planning. Thus, after checking **RG:haste**, PAULINE examines SUPPORT-1 next; it finds the CONVINCE plan appropriate but collects no new sentence topics. Throughout, the states of satisfaction of certain rhetorical goals are updated, in order to enable appropriate restrictive planning.

PRAGS: start topic organization yet? (checking haste and topic derivation)

-->checking rhet goal: RG:HASTE

...no

Topic collector considering new topic: SUPPORT-1

Deciding whether to convince about SUPPORT-1's affect

CONVINCE-PLAN is appropriate

Now applying *CONVINCE-PLAN*'s steps to SUPPORT-1

.....

SUPPORT-1 spawned the following topics: (GOAL-1 SUPPORT-1)

Similarly, GOAL-1 provides no new topics either. Now all the topics off CONSTRUCT-1 have been examined, and, rather than continue collecting topics off the other inputs (which occurs under conditions that call for extremely well planned, complex, unpressured text), the program can start the topic organization phase. Three syntax goals are created, one for each candidate sentence topic:

Topic collector sending to planner:

<Goal to express CONSTRUCT-1

Nomination: Goal CONVINCE-GOAL; Derivation: original topic>

<Goal to express SUPPORT-1

Nomination: Goal CONVINCE-GOAL; Derivation: Parent CONSTRUCT-1; Role: RESULT>

<Goal to express GOAL-1

Nomination: Goal CONVINCE-GOAL; Derivation: Parent CONSTRUCT-1; Role: SUBGOAL-TO>

The topic organization phase includes the following tasks (some of which can be ignored under certain conditions): topic reordering, insertion of explicit opinions, interpretation, phrasal juxtaposition, new topic inclusion. The first task, reordering, is achieved under guidance of a (prescriptive) paragraph construction plan (of which PAULINE has two). In order to open a paragraph, this plan prescribes using the initial topic and the goal it serves:

Starting to organize topics into phrases...

Reordering input topics

PRAGS: is there time to do REORDERING-INPUT? ...yes

Final order is: (CONSTRUCT-1 SUPPORT-1 GOAL-1)

Proceeding to next stage of paragraph...

Grouping together topics suitable for introducing initial topic

Using topics (CONSTRUCT-1 GOAL-1)

PRAGS: is there time to do SLANTING with (CONSTRUCT-1 GOAL-1)? ...yes

Getting affect of CONSTRUCT-1 (slanting strategy is ENHANCE)

Getting affect of GOAL-1 (slanting strategy is NONE)

Starting to build phrases with CONSTRUCT-1

Next, the planner checks whether the rhetorical goals of opinion call for the inclusion of a sentence of explicit opinion (such as "I am angry about Yale's actions"), and if so whether conditions allow the insertion of such a sentence. This is a restrictive task. Since, however, conditions are not favorable (amongst others, no text has been produced yet, and the rhetorical goal **RG:timidity** prevails during monitor conflict resolution), no sentence is included...

Checking active opinion strategies (COMMENT) on CONSTRUCT-1

Rhet goal conflict resolution...

-->checking rhet goal: RG:FORCE RG:VERBOSITY RG:TIMIDITY RG:PARTIALITY

Chosen to satisfy RG:TIMIDITY

No applicable strategy found

Following this, the program checks whether the rhetorical goals have activated any interpretation inferences. As explained in chapter 3, this task is executed both bottom-up and top-down. Three possible interpretations are found bottom-up, but none of them are affectively suitable (remember, the program is speaking as a protester!):

Checking active opinion strategies (INTERPRETATION) on CONSTRUCT-1

Found possible interpretations:

[<CONFIG-TAKE-CONTROL> on CONSTRUCT-1 :

[ACTOR : STUDENTS-1 (= ?X)]

[INSTR : INTO-PLAZA-1]

[TO : STUDENTS-1 (= ?X)]]

Working out *SELF*'s affect for INTERP-TAKE-CONTROL.434: no applicable strategy

[<CONFIG-CONFRONTATION> on CONSTRUCT-1 :

[ACTOR : STUDENTS-1 (= ?X)]

[RELATIONS : SBG:CONS1-G1]

[CONC2 : GOAL-1 (= ?Y)]

[RELATIONS : OPPO:G1-G2]

[CONC2 : GOAL-2 (= ?Z)]

[ACTOR : YALE (= ?W)]

[OPPOSITES: : () (= ?W) (= ?X)]

[RELATIONS : OPPO:G1-G2]

[CONC1 : GOAL-1 (= ?Y)]]

Working out *SELF*'s affect for STUDENTS-1

Working out *SELF*'s affect for INTERP-CONFRONT-CONCEPT.436: no applicable strat

[<CONFIG-PUNISH> on CONSTRUCT-1 :

[ISA : CONSTRUCT (= T) (= ?P)]

[ACTOR : STUDENTS-1 (= ?X)]

[RELATIONS : SBG:CONS1-G1]

[CONC2 : GOAL-1 (= ?Y)]

[RELATIONS : OPPO:G1-G2]

[CONC2 : GOAL-2 (= ?Z)]

[ACTOR : YALE (= ?W)]

[OPPOSITES: : () (= ?W) (= ?X)]

[RELATIONS : OPPO:G1-G2]

[CONC1 : GOAL-1 (= ?Y)]

[RELATIONS : SBG:CONS2-G2]

[CONC2 : GOAL-2 (= ?Z)]

[CONC1 : CONSTRUCT-2 (= T)]

[ISA : CONSTRUCT (= ?P)]]

Working out *SELF*'s affect for INTERP-PUNISH-CONCEPT.438: no applicable strateg

PAULINE checks whether memory contains another instance of a CONSTRUCT to be used as a reminding, and then tries to use a multi-predicate enhancer or mitigator phrase.

Since, however, the hearer doesn't share the program's opinion of GOAL-1, such a phrase may not have the intended effect; thus the program decides not to slant CONSTRUCT-1 in this way.

Starting to work on reminders similar to CONSTRUCT-1

PRAGS: include a reminding? ...yes

Checking if CONSTRUCT-1 is very similar to other topics in memory ... not really

Starting to build affect phrase around CONSTRUCT-1

PRAGS: should use a phrase to ENHANCE the topic?

-->checking rhet goal: RG:PARTIALITY RG:SIMPLICITY

...yes

PRAGS: slant topics on own bias only?

Rhet goal conflict resolution...

-->checking rhet goal: RG:PARTIALITY RG:SIMPLICITY RG:AGGRESSION RG:TIMIDITY

Chosen to satisfy RG:SIMPLICITY

...no

Found no suitable other topic

Finally, though, a strategy does work: the program finds that the relation between its two topics — SUBGOAL-TO — can be expressed as an enhancer by using the linking phrase "as a reminder to". (Since this phrase requires the use of the desired state, not of the goal itself, PAULINE changes the topic to DIVESTED-1).

Starting to cast topic CONSTRUCT-1 into relational phrase

Searching for another topic from

- Goals: (GOAL-1)

PRAGS: is CONSTRUCT-1 simple enough to be conjoined? ...yes

Will use GOAL-1 (a SUBGOAL-TO relation of CONSTRUCT-1) to build a phrase

Checking active opinion strategies (LINK-PHRASE-GOAL) on CONSTRUCT-1

(slanting strategy on CONSTRUCT-1 is ENHANCE)

Running strategy to use phrase <PHRASE %AS-A-REMINDER>

(altering topic from GOAL-1 to DIVESTED-1)

Strategy successful

The program has now assembled the composite syntax goal to build the first sentence, starting with the protesters' goal, and then stating their action:

Finished organizing topic CONSTRUCT-1; built up syntax goals:

(**<Goal to express DIVESTED-1**

Nomination: Goal CONVINCE-GOAL activated plan; role is SUBGOAL-TO

Derivation: Parent CONSTRUCT-1 via relations (GOOD-GOALS *CONVINCE-PLAN*)

Affect: Desired affect is GOOD; Goals call for () strategy

With preceding words (*DOT* AS A REMINDER TO)

Followed by

(**<Goal to express CONSTRUCT-1**

Nomination: Goal CONVINCE-GOAL activated plan; role is INITIAL-TOPIC

Derivation: original topic

Affect: Desired affect is GOOD; Goals call for ENHANCE strategy

With preceding words (*CMA*>>>)

The phase of topic organization is now complete: each syntax goal has enough information associated with each sentence topic to enable realization to begin. The other two input topics, CONSTRUCT-2 and MTRANS-8, are still waiting for topic collection and subsequent organization, and consequently have almost no information associated with them yet. At this point, realization interleaves with planning: the program continues planning to satisfy its rhetorical goals while realizing the syntax goal it has assembled.

The program starts by deciding what type of sentence — a description, an event, or a relation — to build. Complex sentences such as relations and sentences containing nominalized actions satisfy goals of low simplicity and high formality; here, their satisfaction state (restrictively) selects a normal event sentence. The linking phrase is first:

Generator starting to produce text...

Sentence top: setting up DIVESTED-1 with modifiers (GOOD-GOALS COMPL INFINIT)

PRAGS: nominalize DIVESTED-1? ...no

PRAGS: include pre-sentence adverbial clauses?

Rhet goal conflict resolution...

-->checking rhet goal: RG:FORMALITY RG:SIMPLICITY

Not satisfying any of these goals

... building a normal event sentence

Setting up words or phrases linking new sentence to previous one

```

-----> .DOT.
-----> AS
-----> A
-----> REMINDER

```

Next, the program must select an appropriate sentence subject. Of the two candidates, the one most interesting (central in the story), YALE, is selected. Then the rhetorical goals of opinion are checked to see if opinions should be added, and the noun group is built:

Searching for a subject for sentence about DIVESTED-1

PRAGS: which slot to select from (ACTOR OBJECT) as SUBJECT?

Rhet goal conflict resolution...

-->checking rhet goal: RG:HASTE RG:SIMPLICITY RG:FORMALITY

Ordering (ACTOR OBJECT) of DIVESTED-1 ...order is (ACTOR OBJECT)

Found sentence subject: YALE

Building noun group of YALE

Checking active opinion strategies on YALE

Checking whether to include a personal opinion about YALE

PRAGS: should say ng stress on YALE in NG-OPINION?

Rhet goal conflict resolution...

-->checking rhet goal: RG:PARTIALITY RG:AGGRESSION

Working out *UNIV*'s affect for YALE

...no

Searching for head noun of YALE

Checking active opinion strategies (HEAD-NOUN) on YALE

No applicable strategy found

Checking whether to include hearer's opinion about YALE

PRAGS: mention YALE's role for hearer in NG-ROLE?

Rhet goal conflict resolution...

-->checking rhet goal: RG:HEARER-REF RG:AGGRESSION

Chosen to satisfy RG:AGGRESSION

...no

Building up noun group of form (ROLE ART PRE HEAD POST)

```

-----> TO
-----> YALE
-----> UNIVERSITY

```

At this point, a verb must be selected, because verbs determine the nature of the predicate. As described in chapter 6, the memory network is searched for a pragmatically suitable verb; also, aspects of the topic that may be included in the predicate are tested against the rhetorical goals.

Building predicate with DIVESTED-1

```

Searching for a verb to express DIVESTED-1
Searching concept hierarchy (against AFFECT)  DIVESTED-1 DIVESTED
Found verb DIVEST
Filtering aspects of DIVESTED-1 to build sentence predicate
(using (VERB (OBJ RECIP FROM) (POST-SENT ()) (RECIP))))
PRAGS: can COMPANIES-SA, (the RECIP of DIVESTED-1), be said?
PRAGS: does affect rule allow COMPANIES-SA?
      (relating affect of COMPANIES-SA to current AFFECT GOAL)
Working out *SELF*'s affect for COMPANIES-SA
...yes
PRAGS: does interest rule allow COMPANIES-SA?
Examining hearer KNOWLEDGE
-->checking rhet goal:  RG: HASTE
-->checking hearer interest level
-->checking own interest level
...yes
Now planning to say:
(<Goal to expand DIVEST using #{Procedure 63 SAY-VERB}
  and instrs (GOOD-GOALS GOOD-RELATIONS COMPL INFINIT)>
<Goal to expand COMPANIES-SA using #{Procedure 64 SAY-OBJECT}
  and instrs (GOOD-GOALS GOOD-RELATIONS COMPL INFINIT)>
<Goal to expand DIVESTED-1 using #{Procedure 65 SAY-SENT-END}
  and instrs (GOOD-GOALS GOOD-RELATIONS COMPL INFINIT)>>)

```

The syntax goals that realize into the predicate are treated in order:

Computing appropriate stress for DIVESTED-1 itself

Checking active opinion strategies (MITIGATING-HAVE-AUX-VERB) on DIVESTED-1

No applicable strategy found

----> TO

----> DIVEST

Building noun group of COMPANIES-SA

Checking active opinion strategies on COMPANIES-SA

Checking whether to include a personal opinion about COMPANIES-SA

PRAGS should say ng stress on COMPANIES-SA in NG-OPINION?

Working out 'UNIV.'s affect for COMPANIES-SA

no

Searching for head noun of COMPANIES-SA

Checking active opinion strategies (HEAD-NOUN) on COMPANIES-SA

No applicable strategy found

Checking whether to include hearer's opinion about COMPANIES-SA

PRAGS mention COMPANIES-SA's role for hearer in NG-ROLE?

Rhet goal conflict resolution

-->checking rhet goal RG HEARER-REF RG AGGRESSION

Chosen to satisfy RG AGGRESSION

no

Building up noun group of form (ROLE ART PRE HEAD POST)

----> FROM

PRAGS how many of COMPANIES-SA's 1 aspects to say

as PRE-NOUN-MODS?

-->checking rhet goal RG HASTE RG SIMPLICITY RG FORMALITY

will say 1 aspects

Computing appropriate stress for COMPANIES-SA's filler of NUMBER

Checking active opinion strategies (ADJECTIVE-NUMBER) on COMPANIES-SA

No applicable strategy found

----> COMPANIES

----> DOING

----> BUSINESS

----> IN

====> SOUTH
 ====> AFRICA

Next, the run-on sentence is started. As before, PAULINE must check what type of sentence the active rhetorical goals call for...

 Sentence top setting up CONSTRUCT-1 with modifiers (INITIAL-TOPIC ENHANCE COMPL)
 PRAGS: nominalize CONSTRUCT-1?
 PRAGS: include pre-sentence adverbial clauses?
 Rhet goal conflict resolution...
 -->checking rhet goal: RG FORMALITY RG SIMPLICITY
 Not satisfying any of these goals
 -->checking rhet goal: RG FLORIDITY
 ... building a normal event sentence

Setting up words or phrases linking new sentence to previous one

====> *CMA*

 Building noun group of STUDENTS-1
 Checking active opinion strategies on STUDENTS-1

 Searching for a subject for sentence about CONSTRUCT-1
 Found sentence subject: STUDENTS-1
 Checking whether to include a personal opinion of STUDENTS-1
 Searching for head noun of STUDENTS-1
 Checking active opinion strategies (HEAD-NOUN) on STUDENTS-1
 No applicable strategy found
 Building up noun group of form (ROLE ART PRE HEAD POST)
 Checking whether to include hearer's opinion about STUDENTS-1
 PRAGS: mention STUDENTS-1's role for hearer in NG-ROLE?
 PRAGS: can PLUR. (the NUMBER of STUDENTS-1), be said?
 ...will say 1 aspects

While building the noun group, a (restrictive) strategy of opinion the strategy to
 maximize the number of people performing GOOD actions fires:

Checking active opinion strategies (ADJECTIVE-NUMBER) on STUDENTS-1

Running strategy to slant claimed support

Found MAXIMIZE

```

----> A
----> LARGE
----> NUMBER
----> OF
----> STUDENTS

```

Building predicate with CONSTRUCT-1

Searching for a verb to express CONSTRUCT-1

Searching concept hierarchy (against AFFECT): CONSTRUCT-1 CONSTRUCT

Found verb CONSTRUCT

Filtering aspects of CONSTRUCT-1 to build sentence predicate

Now planning to say:

```

(<Goal to expand CONSTRUCT using #{Procedure 63 SAY-VERB}
  and instrs (INITIAL-TOPIC ENHANCE COMPL)>
  <Goal to expand SHANTYTOWN-1 using #{Procedure 64 SAY-OBJECT}
    and instrs (INITIAL-TOPIC ENHANCE COMPL)> (
    <Goal to expand INTO-PLAZA-1 using #{Procedure 66 SAY-INSTR}
      and instrs (INITIAL-TOPIC ENHANCE COMPL)>
    <Goal to expand BEINECKE-PLAZA-1 using #{Procedure 66 SAY-LOC}
      and instrs (INITIAL-TOPIC ENHANCE COMPL)>
    <Goal to expand TIME-1 using #{Procedure 68 SAY-TIME}
      and instrs (INITIAL-TOPIC ENHANCE COMPL)>
    <Goal to expand CONSTRUCT-1 using #{Procedure 65 SAY-SENT-END}
      and instrs (INITIAL-TOPIC ENHANCE COMPL)>))

```

Before saying the verb, the program checks whether any restrictive strategies are applicable; though the strategy to mitigate the verb using "have" could apply, it is not found to be appropriate:

Checking active opinion strategies (MITIGATING-HAVE-AUX-VERB) on CONSTRUCT-1

No applicable strategy found

----> CONSTRUCTED

 Building noun group of SHANTYTOWN-1

Checking active opinion strategies on SHANTYTOWN-1

 Checking whether to include a personal opinion about SHANTYTOWN-1

PRAGS: should say ng stress on SHANTYTOWN-1 in NG-OPINION? ...no

Checking active opinion strategies (HEAD-NOUN) on SHANTYTOWN-1

No applicable strategy found

Checking whether to include hearer's opinion about SHANTYTOWN-1

Building up noun group of form (ROLE ART PRE HEAD POST)

====> A

====> SHANTYTOWN

PRAGS: say (WINNIE MANDELA CITY), (the NAME of SHANTYTOWN-1)? ...yes

====> ---

====> NAMED

====> WINNIE

====> MANDELA

====> CITY

====> ---

====> ON

====> BEINECKE

====> PLAZA

====> IN

====> EARLY

====> APRIL

And so the first sentence ends. At this stage, PAULINE has satisfied, to varying degree, its rhetorical goals; now it gets the opportunity to do so using plans of a more prescriptive nature. More topics must be found. PAULINE still has the two input topics CONSTRUCT-2 and MTRANS-8. As with CONSTRUCT-1, the program finds that the CONVINCE plan is appropriate; when applied to CONSTRUCT-2, it finds the topics GOAL-2, SUPPORT-2 (the local community's outrage at Yale's action), and ARREST-1 (the arrest of the students). For example:

 Topic collector considering new topic: ARREST-1

CONVINCE-PLAN is appropriate

Now applying *CONVINCE-PLAN*'s steps to ARREST-1

Topic ARREST-1 spawned the following topics: (GOAL-2 SUPPORT-2 ARREST-1)

Eventually, when nothing more is found, the topic organization phase is begun. After supervising the reordering of the topics, the paragraph construction plan groups together the topics suited to state and expound a subsequent topic. Then, as before, the presence of rhetorical goals of opinion that suggest the inclusion of sentences with explicit opinions is checked:

Starting to organize topics into phrases...

Reordering input topics

Final order is: (CONSTRUCT-2 ARREST-1 SUPPORT-2 GOAL-1 GOAL-2)

Proceeding to next stage of paragraph...

Grouping together topics suitable for stating and expounding a subsequent topic

-->checking rhet goal: RG:VERBOSITY

Using topics (CONSTRUCT-2 ARREST-1 SUPPORT-2)

PRAGS: is there time to do TOPIC-ORG with CONSTRUCT-2? ...yes

Starting to build phrases with CONSTRUCT-2

Checking active opinion strategies (COMMENT) on CONSTRUCT-2

No applicable strategy found

Following this, PAULINE succeeds in interpreting CONSTRUCT-2 as an instance of someone (Yale) making somebody else (the officials) do their dirty work (the demolishing). This new interpretation is added into memory, and it replaces the original candidate sentence topic:

Checking active opinion strategies (INTERPRET) on CONSTRUCT-2

Found interpretations:

```

[<CONFIG-CAUSE-CONCEPT> on CONSTRUCT-2 :
  [ACTOR : OFFICIALS (= ?X)]
  [RELATIONS : SBG:CONS2-G2]
  [CONC2 : GOAL-2]
  [ACTOR : YALE (= ?Y)]]
Working out *SELF*'s affect for OFFICIALS
Working out *SELF*'s affect for INTERP-CAUSE-CONCEPT.585
Running strategy to interpret as CAUSE-CONCEPT
Strategy successful

```

```

-----
Found that CAUSE-CONCEPT is a new interpretation of CONSTRUCT-2
and related concepts
...building new concept INTERP-CAUSE-CONCEPT.585 about it
Indexing INTERP-CAUSE-CONCEPT.585 under high-level concept CAUSE
Adding back-links from CONSTRUCT-2 to <CONFIG-CAUSE-CONCEPT>
-----

```

```

Starting to work on reminders similar to CONSTRUCT-2
Found no reminders from ()
Checking if CONSTRUCT-2 is very similar to other topics ...not really
Starting to build affect phrase around CONSTRUCT-2
Found no suitable other topic
Starting to cast topic CONSTRUCT-2 into relational phrase
Searching for another topic from
- Goals: (GOAL-2)
- Similar topics: (GOAL-2 ARREST-1 SUPPORT-2)
Will use ARREST-1 (a SIBLING relation of CONSTRUCT-2) to build a phrase
-->checking rhet goal: RG:FORMALITY with phrase <PHRASE %AND1>
Checking active opinion strategies (COMMENT) on CONSTRUCT-2
No applicable strategy found

```

```

Finished organizing topic CONSTRUCT-2; built up syntax goals:
(<Goal to express INTERP-CAUSE-CONCEPT.473
Nomination: Goal CONVINCE-GOAL activated plan; role is INITIAL-TOPIC
Derivation: original topic
Affect: Desired affect is BAD; Goals call for ENHANCE strategy

```

Expression: Phrase/verb HAVE

Followed by

(**<Goal to express ARREST-1**

Nomination: Goal CONVINCE-GOAL activated plan; role is RESULT

Derivation: Parent INTERP-CAUSE-CONCEPT.473 via relations

(GOOD-RELATIONS *CONVINCE-PLAN*)

Affect: Desired affect is BAD; Goals call for ENHANCE strategy

With preceding words (*SEM* ALSO *CMA* AT THAT TIME *CMA*)>>>)

Though, under certain circumstances, realization of this goal could begin immediately, the program proceeds with organizing the next candidate topics...

PRAGS: start realization? (checking haste and topic derivation)

-->checking rhet goal: RG:HASTE

...no

Considering SUPPORT-2

Checking active opinion strategies (COMMENT) on SUPPORT-2

No applicable strategy found

Checking active opinion strategies (INTERPRET) on SUPPORT-2

as INTERP-CONFRONT-CONCEPT.475: no applicable strategy found

Checking active opinion strategies (COMMENT) on SUPPORT-2

No applicable strategy found

Finished organizing topic SUPPORT-2; built up syntax goals:

(**<Goal to express SUPPORT-2**

Nomination: Goal CONVINCE-GOAL activated plan; role is RESULT

Derivation: Parent INTERP-CAUSE-CONCEPT.473 via relations

(GOOD-RELATIONS *CONVINCE-PLAN*)

Affect: Desired affect is GOOD; Goals call for () strategy>)

At this point, the two syntax goals that derived from CONSTRUCT-2 have been completed, and are sent to the realizer. Still waiting to be handled is the goal to generate MTRANS-8 together with whatever topics can be suitably collected.

Generator starting to produce text...

Sentence top: setting up INTERP-CAUSE-CONCEPT.473 with modifiers (ENHANCE)
... building a normal event sentence

Setting up words or phrases linking new sentence to previous one

====> *DOT*

Searching for pre-subject clauses for INTERP-CAUSE-CONCEPT.473

====> AT

====> 5:30

====> AM

====> ON

====> APRIL

====> 14

====> *CMA*

Searching for a subject for sentence about INTERP-CAUSE-CONCEPT.473

====> YALE

Searching for a verb to express INTERP-CAUSE-CONCEPT.473

Using preselected verb HAVE

Building a sentence predicate with HAVE

Filtering aspects of INTERP-CAUSE-CONCEPT.473 to build predicate

-->checking hearer's language ability

-->checking conversational conditions

====> HAD

Sentence top: setting up CONSTRUCT-2 with modifiers (COMPL REPORTED ROOT ENHANCE)
... building a normal event sentence

Setting up words or phrases linking new sentence to previous one

Searching for a subject for sentence about CONSTRUCT-2

Searching for head noun of OFFICIALS

Checking active opinion strategies (HEAD-NOUN) on CONSTRUCT-2

Checking active opinion strategies (ADJECTIVE-NUMBER) on CONSTRUCT-2

====> OFFICIALS

Searching for a verb to express CONSTRUCT-2

Building a sentence predicate with DESTROY

====> DESTROY

Searching for head noun of SHANTYTOWN-1

====> IT

Sentence top: setting up ARREST-1 with modifiers (GOOD-RELATIONS ENHANCE COMPL)

... building a normal event sentence

Setting up words or phrases linking new sentence to previous one

====> *SEM*

====> ALSO

====> *CMA*

====> AT

====> THAT

====> TIME

====> *CMA*

Searching for a subject for sentence about ARREST-1

Found sentence subject: POLICE

====> THE

====> POLICE

Searching for a verb to express ARREST-1

Building a sentence predicate with ARREST

====> ARRESTED

Searching for head noun of STUDENTS-2

====> 76

====> STUDENTS

Sentence top: setting up SUPPORT-2

... building a normal event sentence

Setting up words or phrases linking new sentence to previous one

====> *DOT*

Searching for pre-subject clauses for SUPPORT-2

Searching for a subject for sentence about SUPPORT-2

Checking active opinion strategies (ADJECTIVE-NUMBER) on COMMUNITY-2

Running strategy to slant claimed support

Found MAXIMIZE

====> A

====> LARGE

```

----> NUMBER
----> OF
----> LOCAL
----> POLITICIANS
----> AND
----> FACULTY
----> MEMBERS

```

Searching for a verb to express SUPPORT-2

Searching for a suitable MTRANS verb

Building a sentence predicate with EXPRESS

```

----> EXPRESSED

```

Searching for head noun of CRITICISM

```

----> CRITICISM
----> OF
----> YALE'S
----> ACTION

```

Finally, the program gets to the last topic, MTRANS-8. No new additional topics are found, and hence no interpretations or phrasal juxtapositions occur. Realization is straightforward.

Topic collector considering new topic: GOAL-1

Topic GOAL-1 spawned the following topics: (GOAL-2 GOAL-1)

Starting to organize topics into phrases...

Finished organizing topic MTRANS-8; built up goals:

(<Goal to express MTRANS-8

Nomination: Goal CONVINCE-GOAL activated plan; role is INITIAL-TOPIC

Derivation: original topic

Affect: Desired affect is GOOD; Goals call for () strategy>)

Sentence top: setting up MTRANS-8

... building a norm¹ event sentence

Setting up words or phrases linking new sentence to previous one

```

----> *DOT*
----> FINALLY
----> *CMA*

```

----> THE
 ----> UNIVERSITY
 Searching for a verb to express MTRANS-8
 Searching for a verb to express ALLOW-3
 Building a sentence predicate with GIVE
 ----> GAVE
 ----> THE
 ----> STUDENTS
 ----> PERMISSION
 Searching for a verb to express CONSTRUCT-3
 ----> TO
 ----> REASSEMBLE
 ----> IT

Whole story:

AS A REMINDER TO YALE UNIVERSITY TO DIVEST FROM COMPANIES DOING
 BUSINESS IN SOUTH AFRICA, A LARGE NUMBER OF STUDENTS CONSTRUCTED
 A SHANTYTOWN --- NAMED WINNIE MANDELA CITY --- ON BEINECKE PLAZA
 ON EARLY APRIL. AT 5:30 AM ON APRIL 14, YALE HAD OFFICIALS DESTROY
 IT; ALSO, AT THAT TIME, THE POLICE ARRESTED 76 STUDENTS. A LARGE
 NUMBER OF LOCAL POLITICIANS AND FACULTY MEMBERS EXPRESSED CRITICISM
 OF YALE'S ACTION. FINALLY, THE UNIVERSITY GAVE THE STUDENTS PERMISSION
 TO REASSEMBLE IT.

At the end, PAULINE displays the final satisfaction statuses of the rhetorical goals.
 Note the relatively large number of times the most important goals in this setting --
 partiality, timidity, formality, and aggression -- are satisfied:

Satisfaction of current rhetorical goals:

FORMALITY: 29	DETAIL: 8
FORCE: 10	HASTE: 34

TIMIDITY: 62	PARTIALITY: 36
SPEAKER-REF: 3	HEARER-REF: 14
COLOR: 4	FLORIDITY: 8
RESPECT: 0	SIMPLICITY: 9
WARMTH: 2	VERBOSITY: 13
AGGRESSION: 20	INCITEMENT: 19

Appendix B

A Phrasal Grammar

As described in chapter 6, PAULINE's grammar consists of a set of phrases. They are listed in this appendix.

The phrases can be arranged in a rough hierarchy depending on how much effect they have on the final text. At the level of largest effect, the phrases control the formation of multi-predicate sentences, such as enhancer and mitigator phrases and relations between topics. At the next level, the phrases determine sentence content and organization to form various types of sentences (questions, imperatives). At lower levels, the content and organization of predicates, adverbial clauses, and noun groups are determined.

- PAULINE's phrasal lexicon contains a large number of multi-predicate patterns. When appropriate, depending on the relationships between the sentence topics, the program's planner casts the topics into these patterns. Multi-predicate patterns are used to express the following:

- **Slanting phrases** such as "Not only X, but Y" and "X, however, Y"
- **Reminding phrases** such as "X, which reminds me of Y"
- **Goal-relationship phrases** such as "X in order to Y" and "X so that Y"
- **Result-relationship phrases** such as "X. As a result, Y" and "Y because X"
- **Other relationship phrases** such as "X is larger than Y" and "After X, Y"

- **SAY-SENT-TOP** -- This specialist determines which type of sentence to make and builds the appropriate syntax goal. The decision is based on the input: objects and states are described by **SAY-ATTRIB-SENT**, relations between concepts by **SAY-RELATION-SENT**, and events and state changes by **SAY-EVENT-SENT**. If implemented in **PAULINE**, **SAY-IMPERATIVE** and **SAY-QUESTION** would be included here.
- **SAY-RELATION-SENT** -- Builds a sentence to express the relation between two concepts. The input is the concept representing the relation, which contains a primary (earlier/antecedent/closer) part and a secondary part. This specialist builds one of the sentence patterns

- [**SAY-PRE-SENT** **SAY-COMPL** **SAY-LINK** **SAY-PRE-SENT** **SAY-COMPL**]
 [**SAY-LINK** **SAY-COMPL** , **SAY-COMPL**]

where each **COMPL** specialist is associated with one part of the relation. From the relation concept, the **SAY-LINK** specialist can find a suitable relation word. **PAULINE**'s memory contains the following relations between entities:

1. **Causal:** **CAUSE** or **PRECONDITION** (expressed by "because", "since"); **RESULT** or **ENABLE** (expressed by "because", "enable")
2. **Temporal:** **AFTER** (expressed by "after", "later", "then"); **BEFORE** ("before", "prior to"); **DURING** ("while", "during")
3. **Spatial:** **ABOVE** ("above", "on top of"); **BELOW** ("under", "below", "beneath"); **PROX** ("next to", "beside", "adjoining")
4. **Comparative:** **GREATER** (numerical); **SMALLER**; **MUTUALLY-EXCLUSIVE**
5. **Intergoal:** **SUBGOAL-TO** (an act or a goal serves a goal); **SUBSUMING** (the inverse relation); **OPPOSING** (goals with opposite desires)

- **SAY-ATTRIB-SENT** -- This specialist expands into the list of specialists
 [**PRE-SENT** **SAY-SUBJECT** **SAY-VERB** **SAY-ADVERB** **SAY-ATTRIB**
SAY-POST-SENT]

after selecting from the input and the rhetorical goals which aspect of the input to

describe. If the input is an object, the attribute is an adjective; if an action, an adverb; if a state, a degree; if a state-change, an adverb or a degree. The attribute can be said in various ways:

- "the bag is red"
- "the color is red"
- "the bag's color is red"
- "the color of the bag is red"
- "red is the color of the bag"
- "the bag has a color"
- "the bag has a red color"

- SAY-EVENT-SENT — This specialist expands into the sequence

- [SAY-PRE-SENT SAY-SUBJECT SAY-PREDICATE]

after finding which aspect of the input to make the sentence subject.

To choose the subject, the specialist queries the activated rhetorical goal strategies to find the pragmatically most interesting aspect of the input. If, for example, the goal **RG:haste** has the value *pressured*, no time is wasted on evaluating the affects of various candidates; the **ACTOR** (or, failing that, the **INSTR**) aspect is selected directly. Otherwise, the candidates for sentence subject are ordered by affect (sympathetic to the hearer, when not being *aggressive* or *inciting*); by their relation to the central topic; and by the amount of information represented for each one. Strategies of focus, such as described in [McKeown 82] and discussed in chapter 6, have not been implemented in PAULINE.

The choice of the subject is also related to the verb. Some generators always select the verb first, at this stage of the realization process (see, for example, [Danlos 84]), otherwise, the generator may sometimes produce bad text: for example, some pre-sentence clauses may be prohibited by the verb (still to be chosen) or may get from it a non-standard preposition. Thus, verb choice must be able to inspect the exact words that precede it. In addition, of course, verbs have no way of indicating which

aspect of the input they prefer as subject. PAULINE usually chooses the verb only when it builds the predicate, because this corresponds to the way people usually speak: we often only choose a verb at 'verb time', after the subject has been said, and any idiosyncratic constraints a verb may have either disqualifies the verb or causes a re-start of the sentence. (Why else does the longest intrasentential pause occur just before the verb? Why else do we so often start the sentence again, using a different subject, at that point?) But, of course, we are able to choose the verb before. (Sometimes we have to. For example, in French, the only way to say "Pete misses Mary" is "Mary manque à Pete"; this selection of a non-ACTOR as subject is required by "manquer". However, PAULINE can choose a verb before starting realization (it does so after the interpretation of topics, as described in chapter 3, for example); in such cases, the only action taken is a check of the features of the verb in order to ensure a valid subject has been chosen (otherwise, for example, if the verb "beat" has been chosen, and the ACTOR is Kennedy, then the simplistic strategy of choosing the actor would produce "Kennedy beat Kennedy"). These and similar arguments against straightforward left-to-right generation are made in [Danlos 86].

- SAY-COMPL --- Builds a sentence without the SAY-PRE-SENT specialist, for use in cases such as "(He said that) they went to New York"; i.e., expands into
 - [SAY-SUBJECT SAY-PREDICATE]
- SAY-REL-CLAUSE -- Builds a relative clause. The input marks which aspect is shared by the surrounding syntactic environment; this aspect is associated with the SAY-REL-PRONOUN specialist, and the rest is treated like a sentence
- SAY-PREDICATE --- This specialist builds a sentence predicate from its input. Unless a verb has already been chosen by rhetorical planning, it selects a verb. To get a predicate pattern, it checks the verb; if no idiosyncratic pattern is found, the standard pattern is used:
 - [SAY-VERB SAY-OBJECT SAY-ADVERB SAY-POST-SENT]

In the lexicon, the formative pattern associated with a verb is a list of units. Each unit gives the position of its corresponding entity in the predicate. The absence of

an aspect in the pattern means that the aspect cannot be said; a required aspect is marked mandatory. This information is used for the *inclusion* decisions. The *ordering* is given in the pattern; when various orders are possible, the pattern itself is written as a specialist function that queries the rhetorical strategies for assistance (for example, the typical adverbial clauses of time, instrument, and location are handled by SAY-POST-SENT). The *casting* function is done by associating specialists and aspects of the input, as prescribed by each unit in the pattern. A unit can consist of singles, pairs, or triplets:

- A single element (or the first one of a pair or triplet) is a keyword that indicates which specialist is to provide the syntactic environment of that entry. If the element is not the name of a specialist function, it is taken to be a literal — part of a frozen phrase — which must be said.
- The second element of a pair or triplet indicates which aspect of the input is to be used by the activated specialist function. For example, the pair (SAY-OBJECT object) indicates that the filler of the OBJECT aspect of the input is to be generated as an accusative case noun phrase.
- The third element of a triplet is the preposition to be used in a preposition group. For example, (SAY-LOC to to) indicates that the specialist SAY-LOC is to use the TO aspect filler with the preposition "to".
- The keyword SAY-POST-SENT indicates all the adverbial clauses that can normally be said in an English predicate, each of which has a specialist (SAY-INSTR, SAY-TIME, SAY-LOC, SAY-TO, SAY-FROM, etc.).
- The presence of a literal in the predicate pattern indicates that it must be said. A specialist function keyword indicates its position if it can be included (semantically, if its aspect appears in the input, and pragmatically, if the rhetorical goals allow). Some predicate patterns require the presence of a syntactic environment; this is indicated by a fourth entry in the unit (otherwise, under pragmatic guidance, the generator may produce sentences such as "Kennedy beat in the election"). Sometimes patterns contain parts that may only be said if they dif-

fer from the subject, these parts are also appropriately marked (otherwise, the generator may produce "Pete gave the book to John from Pete").

For example, to express the Conceptual Dependency representation primitive MTRANS, PAULINE has more than 20 words. (MTRANS stands for transfer of information; the aspect OBJECT contains the message and the aspect TO the hearer; see [Schank 72, 75]). Some of these words are "tell" (two versions) and "say":

– "tell":

[SAY-VERB [SAY-OBJECT (aspect TO)] SAY-ADVERB SAY-POST-SENT
that [SAY-COMPL (aspect OBJECT)]]

"He told her quietly yesterday that [she should see the film]"

– "tell-1":

[SAY-VERB [SAY-OBJECT (aspect TO)] SAY-ADVERB SAY-POST-SENT
[SAY-PRED (aspect OBJECT) infinitive]]

"He told her quietly yesterday [to see the film]"

– "say":

[SAY-VERB [SAY-OBJECT (aspect TO) to] SAY-POST-SENT that
[SAY-COMPL (aspect OBJECT)]]

"He said to her yesterday that [she should see the film]"

As discussed before, the verb is chosen when SAY-PREDICATE builds the predicate. Linking into the lexicon from the type of the input concept, PAULINE searches the memory hierarchy until it finds a WORD aspect, which either directly indexes words or provides a discrimination net which indexes other concepts or words. At the top of the hierarchy, elaborate discrimination nets are associated with the conceptual dependency and other primitives to ensure that some verb will always be found.

- SAY-LINK-WORD This specialist controls the use of words said at the beginning of a sentence to link it to the previous sentence, for example "and", "but", "as a result", "so". The question of when to include SAY-LINK-WORD in the expansion stream illustrates the general problem of where to plan how much of the text.

One argument calls for including SAY-LINK-WORD at a relatively high level — say, before SAY-SENT-TOP. The rationale is that these functions, on or just below the level of the planner, are the only ones with the requisite breadth of view to decide on proper link words, and are therefore in the position to plan for them. For example, in JUDGE example, if the top-level say-function decides to say the final result after the fight itself, it can plan on saying “FINAL-ACTION and finally FINAL-RESULT”. There is a problem with this: sometimes FINAL-ACTION and FINAL-RESULT are too different to compare comfortably with “and”, for example in “Sam hit Jim. His action was justified, and finally Jim died.” The top-level function would have to pre-plan all the way down to the justification to handle this.

The other alternative is to make SAY-LINK-WORD much more intelligent. It must be able to compare the previous and current sentences and decide what link words are appropriate. The simplest version would just say “and” for everything — much as small children do — and a more sophisticated one may still produce sentences such as “She liked the lawn and gardening”. This sentence is an example of faulty parallel construction, a topic that receives much attention in stylistic handbooks (this example is from [Baker 66, p 106]). The fact that people require explicit training in this matter suggests that the decision to make SAY-LINK-WORD do the work is the right one.

PAULINE's construction of an argument and its indication of the relations among parts leaves much to be desired. It does not construct some type of argument graph, as discussed in chapter 6 and in [Birnbaum 85], from which this specialist can choose linking words with which to preface each sentence. In fact, by the time that SAY-LINK-WORD is expanded, most records of the derivation of the situation have vanished and the specialist has to discriminate among its options by whatever residues are left. PAULINE adds some topic collection information (typically, what relation the sentence topic bears to its ancestor in the collection process, as used in the topic collection plans described in chapter 6) to the syntax goal, to enable SAY-LINK-WORD to use phrases such as “one result is” and “a good example is”. The work on rhetorical structure theory ([Mann & Thompson 85]) can be extended and used by

this specialist.

- **SAY-PRE-SENT-CLAUSES** -- Returns the specialists of all the pre-subject clauses. Check pragmatic strategies to see how many aspects to include, and how many to include before the subject. (As described in chapter 5, the values *complex* for the rhetorical goal **RG:simplicity** and *highfalutin* for **RG:formality** require many clauses in this position)
- **SAY-PREPGROUP** -- Returns [preposition SAY-NOUN-GROUP]
- **SAY-SUBJECT, SAY-OBJECT, SAY-POSSESSIVE** — Set the case of the syntax goal to *nominative*, *accusative*, and *genitive*, respectively, and returns SAY-NOUN-GROUP
- **SAY-ACT-AS-OBJECT** -- Builds a noun group of an event or state change (for example, "Sam's shot" if the input has been said before; "John's being shot by Sam", passive for pro-victim affects; "Sam's shooting John" otherwise)
- **SAY-NOUN-GROUP** — Decides whether to pronominalize or not; if not, selects a head noun and returns
 - [SAY-ART SAY-PRE-NOUN-MODS SAY-HEAD-NOUN SAY-POST-NOUN-MODS]
 Sets up the context for the noun group: number, case, gender (the latter are used in languages with case and gender declension), etc.

To find a suitable head noun for the input, PAULINE executes a series of tests. If it has just been speaking about the input topic then it can simply say its name or word. It checks the rhetorical goals whether an opinion should be chosen as the head noun (producing "that jerk, Kennedy"). If not, from all the aspects of the input, it excludes those not suitable for the head noun, those that have no defined word, and those that have been said, and then orders them by the rhetorical strategies and uses the preferred aspect. If pragmatics produce no preferences, it uses a predefined default ordering.

PAULINE has four pronominalization strategies:

1. A simple *most-recent* rule, based on the fact that, in English, pronouns carry number and gender information: when PAULINE says a representation element, it stores a triplet — the input, its number, its gender — on a list. The next time PAULINE references an element, it checks whether the most recent entry on the reference list with the same number and gender is the same input element; if so, it pronominalizes. This simple-minded strategy is often described in grammar books. Of course, it doesn't always work: "Pete and John went to the shop. He came back first" — where he is John; or, better: "Pete saw Mike in the store. He..." The "he", according to this strategy, refers to Mike. But in "Pete saw Mike in the store. He told him about Mary", people quite naturally assume subject-subject and object-object correspondence.
2. The obvious *case-correspondence* strategy: if the current sentence subject (object) is the same as the previous sentence subject (object), pronominalize. Then in "Pete saw Mike in the store. He..." "he" refers to Pete. For this strategy, PAULINE obviously also stores case information on the reference list.
3. The conjunction of the above two strategies. This is stricter than either of them: it only allows pronominalization when the input entity corresponds to the most recent match of number, gender, *and* case.
4. The disjunction of the above two strategies. This is less strict than either of them. This strategy gives the most natural text, since it denies pronouns *only* in cases such as "Pete saw Mike in the store. He told Pete about Mary", where the subject and object swap *and* they have the same number and gender.

The most recent strategy is clearly inadequate, and the corresponding-case is better but not yet satisfactory. A better strategy would take into account something like the notion of focus of attention discussed in chapter 6, as worked out by Grosz [Grosz 77], Sidner [Sidner 79], and McKeown [McKeown 82], which allows you to pronominalize the input if it is the current focus of attention and has been said. However, to do proper pronominalization, you require unlimited use of the whole inference capabil-

ity of the system. And that is a big engine to run every time you want to decide whether to say a pronoun or not (see [Appelt 86] for a serious attempt at doing this). Though people obviously sometimes do it, we probably often use some shortcut strategy instead. The disjunction strategy is one such shortcut. Since people usually pronominalize as much as possible in normal circumstances, this strategy produces the most natural text.

- The following build parts of noun groups: SAY-ART, SAY-HEAD-NOUN, SAY-PRE-NOUN-MODS, SAY-POST-NOUN-MODS, SAY-PRONOUN, SAY-REL-PRONOUN. Pre- and post-nominal modifiers are selected by the rhetorical strategies from the aspects of the input, and ordered by their affective preference. If none exists, a default order is used. Some modifiers can only appear before or after the head noun; of the rest, PAULINE places equally many in each position, unless required to do otherwise by the strategies described in chapters 3 to 5.
- SAY-VERB -- Selects the appropriate tense form and auxiliary verbs and conjugates the selected verb
- SAY-PRE-VERB-MODS -- Modify the verb by saying either of [INTENT ALSO], as in "Mike also hit Jim"; if the action has been represented as intentional, say so; otherwise, STRESS can be said
- SAY-ADVERB -- Modify the verb by saying any or all of
 [PREP-MODIFYING-VERB STRESS ADVERB REPEATER BACK/AGAIN]
 as in "Jim knocked Mike down very hard again"; for example:
 - INTENT: "purposely", "intentionally", "wilfully", "on purpose", /
 "accidentally", "unintentionally", "by accident"
 - ALSO: "also"
 - STRESS: "really", "easily", "badly" / "only", "just", "merely", "narrowly"
 - PREPOSITION: any preposition associated with the verb
 - REPEATER: "repeatedly" / "once"

- BACK/AGAIN: "back" / "again"
- ADVERB: any adverb

The prepositions "down" and "into" in the sentences

- "Jim knocked Mike down again"
- "Jim bumped into Mike again"

derive from different sources: "down" modifies the verb "knock" — Jim could certainly have knocked Mike over, or to the floor, or under the table. However, "bump into" is a frozen phrase — you cannot change the preposition to "bump to" or "bump over" and retain the meaning. This difference is reflected in the generator's representation for the two cases. In the former, the specialist SAY-ADVERB finds the semantic equivalent of "down" in the input representation and says "down"; it could just as easily have found "up" or "over". In the latter, the "into" is part of the sentence form for the verb "bump into", and the OBJECT must be said with the preposition.

Most of the specialists listed above correspond to traditional phrase structure symbols such as NOUN GROUP and PREDICATE. But, just as specialists may be used to build multi-predicate sentences, they may be used for other purposes. For example, people talk about money in various highly idiomatic ways. All PAULINE's phrasal knowledge relating to money is grouped together in a single specialist. The phrasal lexicon contains the following specialists:

- SAY-MONEY — Realize the phrases "the measly \$15" (as head noun); "the 95c book" (as pre-nominal modifier); and "a green truck *worth 900 bucks*" (as post-nominal modifier). As with all the other variations, options are selected by referring to the relevant pragmatic criteria: for example, "bucks" is not selected when pragmatics call for being very formal.
- The specialists SAY-AMOUNT and SAY-DIFFERENCE express other numerical amounts. The former builds the patterns
 "a {size} number/amount (of {unit})"

"[amount] [unit]"

- "the most/fewest/X number of [unit]"

the latter expresses the numerical difference between two amounts.

- The following are adverbial clause specialists: SAY-SOURCE, SAY-RECIPIENT, SAY-TO, SAY-FROM, SAY-TIME, SAY-LOC, SAY-INSTRUMENT, SAY-PRE-INSTR, SAY-MEASURE. Each specialist is able to produce various English forms. For example, depending on the nature of the input it receives, SAY-TIME can produce:
 - "now / today / tomorrow / yesterday": input is a predefined day concept
 - "at 5 o'clock today / yesterday": input is a MEASURE concept
 - "in the future / past": input is the concept FUTURE or PAST
 - "15 hours from now": input is a MEASURE, WHEN aspect is FUTURE or PAST
 - "at 3 o'clock": input is an INSTANCE
 - The following specialists are all used to build noun groups: SAY-NAME, SAY-AGE, SAY-GENDER, SAY-NUMBER, SAY-TITLE, SAY-NATION, SAY-RESID, SAY-WEARING, SAY-OCCUP, SAY-DESCRIP, SAY-SIZE, SAY-COLOR, SAY-OWNER, SAY-OPINION, SAY-ROLE, SAY-MONEY. Each specialist must be able to produce various forms, depending on its position. For example, SAY-AGE can produce the following forms:
 - *predicate*: "... is 29 years old"
 - *head-noun*: "the fat but pretty 29-year-old from Irkutsk"
 - *pre-head-noun*: "the 29-year-old fat woman, Marta"
 - *post-head-noun*: "Marta, a 29-year-old, ..."
 and SAY-GENDER the following:
 - *predicate*: "... is female"
 - *head-noun*: "the green man from Mars"
 - *pre-head-noun*: "the large male dinosaur"
 - *post-head-noun*: "the student who is male..."
- The inclusion and ordering decisions of topic aspects are all made by the rhetorical strategies, as described in chapters 3, 4, and 5.

Bibliography

1. Abelson, R.P. & Rosenberg, M.J.,
Symbolic Psycho-Logic: A Model of Attitudinal Cognition, in *Behavioral Science*, vol 4, 1958.
2. Allen, J.F. & Perrault, C.R.,
Analyzing Intentions in Dialogues, University of Rochester Technical Report no 150, 1979.
3. Appelt, D.E.,
Planning Natural Language Utterances to Satisfy Multiple Goals, Ph.D. dissertation, Stanford University, 1981.
4. Appelt, D.E.,
Planning Natural-Language Utterances, AAAI Conference Proceedings, 1982.
5. Appelt, D.E.,
Planning English Sentences, Cambridge: Cambridge University Press, 1985.
6. Appelt, D.E.,
Towards a Plan-Based Theory of Referring Actions, in **Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics**, Kempen G. (ed), Boston: Kluwer Academic Publishers, 1987.
7. Atkinson, J.M.,
Understanding Formality: the Categorization and Production of 'Formal' Interaction, in *British Journal of Sociology*, vol 33 no 1, 1982.
8. Bain, W.M.,
Toward a Model of Subjective Interpretation, Yale University Technical Report no 324, 1984.

9. Bain, W.M.,
Case-based Reasoning: A Computer Model of Subjective Assessment, Ph.D. dissertation, Yale University, 1985.
10. Bain, W.M.,
A Case-based Reasoning System for Subjective Assessment, AAAI Proceedings, 1986.
11. Baker, S.,
The Complete Stylist, New York: Crowell, 1966.
12. Becker, J.D.,
The Phrasal Lexicon, Bolt, Beranek and Newman Technical Report no 3081, 1975.
13. Bienkowski, M.A.,
A Computational Model for Extemporaneous Elaborations, Princeton University Cognitive Science Laboratory Technical Report no 1, 1986.
14. Birk, N.P. & Birk, G.B.,
Understanding and Using English, fourth edition, New York: Odyssey Press, 1965.
15. Birnbaum, L.A.,
A Functional Approach to the Representation of Arguments, in **Advances in Cognitive Science**, Sharkey, N. (ed), Chichester: Ellis Horwood, 1985.
16. Birnbaum, L.A.,
Integrated Processing in Planning and Understanding, Ph.D. dissertation, Yale University, 1986.
17. Birnbaum, L.A., Flowers, M. & McGuire, R.,
Towards an AI Model of Argumentation, AAAI Conference Proceedings, 1980.
18. Bloomfield, L.,
An Introduction to the Study of Language, New York: Holt, 1914.

19. Bobrow, D.G. & Winograd, T.,
An Overview of KRL, a Knowledge-Representation Language, in *Cognitive Science*,
 vol 1 no 1, 1977.
20. Bock, J.K.,
*Toward a Cognitive Psychology of Syntax: Information Processing Contributions to
 Sentence Formulation*, in *Psychological Review* no 89, 1982.
21. Bock, J.K.,
Exploring Levels of Processing in Sentence Production, in **Natural Language Gen-
 eration: Recent Advances in Artificial Intelligence, Psychology, and Lin-
 guistics**, Kempen G. (ed), Boston: Kluwer Academic Publishers, 1987.
22. Brown, P. & Levinson, S.C.,
Universals in Language Usage: Politeness Phenomena, in **Questions and Polite-
 ness: Strategies in Social Interaction**, Goody E. (ed), Cambridge: Cambridge
 University Press, 1978.
23. Bruce, B., Collins, A., Rubin, A.D. & Gentner, D.,
A Cognitive Science Approach to Writing, Bolt, Beranek and Newman Technical
 Report no 89, 1978.
24. Bühler, K.,
Sprachtheorie, Jena: Fischer, 1934.
25. Carbonell, J.G.,
Intentionality and Human Conversations, TINLAP-2 Conference Proceedings, Uni-
 versity of Illinois (Urbana), 1978.
26. Carnap, R.,
Foundations of Logic and Mathematics, in **International Encyclopedia of Unified
 Science**, vol. 1, Neurath O., Carnap R. & Morris C.W. (eds), 1938.
27. Carnap, R.,
Meaning and Necessity, Chicago: University of Chicago Press, 1956.

28. Carnap, R.,
Introduction to Semantics, Cambridge: Harvard University Press, 1959.
29. Charniak, E. & McDermott, D.V.,
Introduction to Artificial Intelligence, New York: Addison Wesley, 1985.
30. Charniak, E., Riesbeck, C.K. & McDermott, D.V.,
Artificial Intelligence Programming Hillsdale: Lawrence Erlbaum Associates, 1980.
31. Chomsky, N.,
Syntactic Structures, The Hague: Mouton, 1957.
32. Chomsky, N.,
Aspects of the Theory of Syntax, Cambridge: MIT Press, 1965.
33. Chomsky, N.,
Reflections on Language, Cambridge: MIT Press, 1976.
34. Clark, H.H. & Carlson, T.B.,
Context for Comprehension, in **Attention and Performance IX**, Long J. & Baddeley A. (eds), Hillsdale: Lawrence Erlbaum Associates, 1981.
35. Clark, H.H. & Murphy, G.L.,
Audience Design in Meaning and Reference, in **Language and Comprehension**, Le Ny J.F. & Kintsch W. (eds), Amsterdam: North Holland Publishing Company, 1982.
36. Clark, H.H. & Schunk, D.,
Polite Responses to Polite Requests, in *Cognition* no 8, 1980.
37. Clippinger, J.H.,
A Discourse Speaking Program as a Preliminary Theory of Discourse Behavior and a Limited Theory of Psychoanalytic Discourse, Ph.D. dissertation, University of Pennsylvania, 1974.

38. Clippinger, J.H.,
Speaking with Many Tongues: Some Problems in Modeling Speakers of Actual Discourse, TINLAP-1 Conference Proceedings, 1975.
39. Cohen, P.R.,
On Knowing What to Say: Planning Speech Acts, Ph.D. dissertation, University of Toronto, 1978.
40. Cole, P. (ed),
Syntax and Semantics Vol 9: Pragmatics, New York: Academic Press, 1978.
41. Cohen, P.R. & Perrault, C.R.,
Elements of a Plan-Based Theory of Speech Acts, in *Cognitive Science*, vol 3 no 3, 1979.
42. Cowan, G. & McPherson, E.,
Plain English Rhetoric and Reader, New York: Random House, 1977.
43. Danks, J.H.,
Producing Ideas and Sentences, in **Sentence Production: Developments in Research and Theory**, Rosenberg S. (ed), Hillsdale: Lawrence Erlbaum Associates, 1977.
44. Danks, L.,
Some Issues in Generation from a Semantic Representation, IJCAI Conference Proceedings, 1983.
45. Danks, L.,
Génération Automatique de Texts en Langues Naturelles, Ph.D. dissertation, University of Paris, Paris: Masson, 1985.
46. Danks, L.,
A Robust Syntactic Component in Generation, in **Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics**, Kempen G. (ed), Boston: Kluwer Academic Publishers, 1987.

47. Davey, A.,
Discourse Production, Edinburgh: Edinburgh University Press, 1979.
48. De Beaugrande, R.,
Text Production, *Advances in Discourse Processes*, vol XI, ABLEX Publishing Corporation, 1984.
49. De Smedt, K. & Kempen, G.,
Incremental Sentence Production, in **Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics**, Kempen G. (ed), Boston: Kluwer Academic Publishers, 1987.
50. Doyle, R.J., Atkinson, D.J. & Doshi, R.S.,
Generating Perception Requests and Expectations to Verify the Execution of Plans, AAAI Conference Proceedings, 1986.
51. Durfee, E.H. & Lesser, V.R.,
Incremental Planning to Control a Blackboard-Based Problem Solver, Proceedings of the Eighth Conference of the Cognitive Science Society, 1986.
52. Fawcett, R.P.,
Cognitive Linguistics and Social Interaction, Heidelberg: Julius Groos Verlag, 1980.
53. Fikes, R.E., Hart, P.E. & Nilsson, N.J.,
New Directions in Robot Problem Solving, in *Machine Intelligence* no 7, 1972.
54. Fillmore, C.J.,
The Case for Case, in **Universals in Linguistic Theory**, Bach E. & Harms R. (eds), New York: Holt, Reinhart & Winston, 1968.
55. Fillmore, C.J.,
Types of Lexical Information, in **Semantics: an Interdisciplinary Reader in Philosophy, Linguistics, and Psychology**, Steinberg D.D. & Jakobovits L.A. (eds), Cambridge: Cambridge University Press, 1971.

56. Freud, S.,
A General Introduction to Psychoanalysis. Riviere J. (trans). New York: Liveright, 1935.
57. Gasser, M. & Dyer, M.G.,
Speak of the Devil: Representing Deictic and Speech Act Knowledge in an Integrated Lexical Memory. Proceedings of the Eighth Conference of the Cognitive Science Society, 1986.
58. Gazdar, G.,
Pragmatics: Implicature, Presupposition, and Logical Form. New York: Academic Press, 1979.
59. Gazdar, G.,
Pragmatic Constraints on Linguistic Production. in **Language Production**, vol 1, Butterworth B. (ed), New York: Academic Press, 1980.
60. Gibbs, R.,
Contextual Effects in Understanding Indirect Requests. in *Discourse Processes* no 2, 1979.
61. Gibbs, R.,
Your Wish is My Command: Convention and Context in Interpreting Indirect Requests. in *Journal of Verbal Learning and Verbal Behavior* no 20, 1981.
62. Goldman, N.M.,
Sentence Paraphrasing from a Conceptual Base. Stanford University Technical Report, 1973.
63. Goldman, N.M.,
Conceptual Generation. in **Conceptual Information Processing**, Schank R.C. (ed), Amsterdam: North-Holland Publishing Company, 1975.

64. Goody, E. (ed).
Questions and Politeness: Strategies in Social Interaction. Cambridge: Cambridge University Press, 1978.
65. Gordon, D. & Lakoff, G.,
Conversational Postulates, in **Syntax and Semantics 9: Pragmatics**, Cole & Morgan (eds), New York: Academic Press, 1975.
66. Gregory, M.,
Towards 'Communication' Linguistics: a Framework, in **Systemic Perspectives on Discourse**, vol 1, Benson J.D. & Greaves W.S. (eds), 1982.
67. Grice, H.P.,
Meaning, in **Semantics: an Interdisciplinary Reader in Philosophy, Linguistics, and Psychology**, Steinberg D.D. & Jacobovits L.A. (eds), Cambridge: Cambridge University Press, 1971.
68. Grice, H.P.,
Logic and Conversation, in **The Logic of Grammar**, Davidson D. & Harman G. (eds), Dickinson Publishing Company, 1975.
69. Grimes, J.E.,
The Thread of Discourse, The Hague: Mouton, 1975.
70. Gross, M.,
On the Failure of Generative Grammar, in *Language*, no 55, 1979.
71. Gross, M.,
Lexicon-Grammar and the Syntactic Analysis of French, Proceedings of the Computational Linguistics Society Conference Coling 84, 1984.
72. Grosz, B.J.,
The Representation and Use of Focus in Dialogue Understanding, Stanford Research Institute Technical Report no 151, 1977.

73. Grosz, B.J.,
Focusing and Description in Natural Language Dialogs, in **Elements of Discourse Understanding**: Proceedings of a Workshop on Computational Aspects of Linguistic Structure and Discourse Setting, Joshi A. et al. (eds), Cambridge: Cambridge University Press, 1980.
74. Grosz, B.J.,
A Theory of Discourse Structure, Proceedings of the Eighth Conference of the Cognitive Science Society, 1986.
75. Grosz, B.J. & Sidner, C.L.,
Discourse Structure and the Proper Treatment of Interruptions, IJCAI Conference Proceedings, 1985.
76. Halliday, M.A.K.,
Halliday: System and Function in Language, (selected papers, edited by G.R. Kress), Oxford: Oxford University Press, 1976.
77. Halliday, M.A.K.,
Language as Social Semiotic, University Park Press, 1978.
78. Harada, S.I.,
Honorifics, in **Syntax and Semantics 5: Japanese Generative Grammar**, Shibatani M. (ed), New York: Academic Press, 1976.
79. Herrmann, T. & Laucht, M.,
Planung von Äusserungen als Selektion von Komponenten Implikativer Propositionsstrukturen, Universität Mannheim Technical Report no 3, 1978.
80. Hill, A.S.,
The Foundations of Rhetoric, New York: Harper & Brothers, 1892.
81. Horn, L.,
Toward a New Taxonomy for Pragmatic Inference: Q-Based and R-Based Implicature, unpublished manuscript, Yale University, 1984.

82. Householder, F.W.,
Linguistic Speculations, Cambridge: Cambridge University Press, 1971.
83. Hovy, E.H. & Schank, R.C.,
Language Generation by Computer, in **Natural Language Processing**, Bara, B.G. & Guida, G. (eds), Amsterdam: North-Holland Publishing Company, 1984.
84. Hovy, E.H.,
Integrating Text Planning and Production in Generation, IJCAI Conference Proceedings, 1985.
85. Hovy, E.H., 1986a,
Some Pragmatic Decision Criteria in Generation, in **Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics**, Kempen G. (ed), Boston: Kluwer Academic Publishers, 1987.
86. Hovy, E.H., 1986b,
Putting Affect into Text, Proceedings of the Eighth Conference of the Cognitive Science Society, 1986.
87. Hovy, E.H., 1987a,
Two Types of Planning in Language Generation, submitted to 1987 IJCAI Conference.
88. Hovy, E.H., 1987b,
Grammar and a Phrasal Lexicon, in **Papers in Natural Language Generation**, McDonald, D.D. & Bolc, (eds), New York: Springer-Verlag, in press.
89. Hovy, E.H., 1987c,
Generating Natural Language under Pragmatic Constraints, in *Journal of Pragmatics*, vol XI no 6, 1987, forthcoming.
90. Hovy, E.H., 1987d,
Interpretation in Generation, in Proceedings of the 25th Annual Meeting of the Association for Computational Linguistics, 1987, forthcoming.

91. Hovy, E.H., 1987e,
What Makes Language Formal?, submitted to 1987 IJCAI Conference of the Cognitive Science Society.
92. Irvine, J.T.,
Formality and Informality of Speech Events, in *American Anthropologist*, vol 81 no 4. 1979.
93. Jackendoff, R.,
On Katz's Autonomous Semantics, in *Language*, vol 57 no 2, 1981.
94. Jackendoff, R.,
Semantics and Cognition, Cambridge: MIT Press, 1985.
95. Jacobs, P.S.,
PHRED: a Generator for Natural Language Interfaces, University of California (Berkeley) Technical Report no CSD 85/198, 1985.
96. Jacobs, P.S.,
A Knowledge-Based Approach to Language Production, Ph.D. dissertation, University of California (Berkeley), 1985.
97. Jakobson, R.,
Linguistics and Poetics, in **Style in Language**, Sebeok T. (ed), Cambridge: MIT Press, 1960.
98. Jameson, A.,
How to Appear to Be Conforming to the 'Maxims' Even If You Prefer to Violate Them, in **Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics**, Kempen G. (ed), Boston: Kluwer Academic Publishers, 1987.
99. Johnson, P.N. & Robertson, S.P.,
MAGPIE: A Goal-Based Model of Conversation, Yale University Technical Report no 206, 1981.

100. Joncas, E.,
Action Expectations in Social Situations, Ph.D. dissertation, Yale University, 1972.
101. Joshi, A.K.,
Tree Adjoining Grammars and their Relevance to Generation, in **Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics**, Kempen G. (ed), Boston: Kluwer Academic Publishers, 1987.
102. Kaplan, R.M. & Bresnan, J.,
The Mental Representation of Grammatical Relations, Cambridge: MIT Press, 1983.
103. Katz, J.J.,
Propositional Structure and Illocutionary Force, New York: Crowell, 1977.
104. Katz, J.J.,
Chomsky on Meaning, in *Language*, vol 56 no 1, 1980.
105. Kay, M.,
Functional Grammar, Proceedings of the Fifth Annual Meeting of the Berkely Linguistics Society, 1979.
106. Kay, M.,
Functional Unification Grammar: a Formalism for Machine Translation, in *Proceedings of the Tenth International Conference on Computational Linguistics*, Stanford University, 1984.
107. Keenan, J.M., MacWhinney, B. & Mayhew, D.,
Pragmatics in Memory: A Study of Natural Conversation, in **Memory Observed**, Neisser, U. (ed), San Francisco: William Freeman Press, 1982.
108. Kempen, G.,
Directions for Building a Sentence Generator which is Psychologically Plausible, unpublished paper, Yale University, 1976.

109. Kempen, G.,
Conceptualizing and Formulating in Sentence Production, in **Sentence Production: Developments in Research and Theory**, Rosenberg S. (ed), Hillsdale: Lawrence Erlbaum Associates, 1977.
110. Kempen, G. & Hoenkamp, E.,
A Procedural Grammar for Sentence Production, University of Nijmegen, The Netherlands, 1978.
111. Kroch, A. & Joshi, A.K.,
Analyzing Extraposition in a Tree Adjoining Grammar, in **Syntax and Semantics (Discontinuous Constituents)**, Huck G. & Ojeda A. (eds), New York: Academic Press, 1986.
112. Kukich, K., Ph.D. dissertation, University Of Pennsylvania, 1983.
113. Kuno, S.,
The Structure of the Japanese Language, Cambridge: Harvard University Press, 1973.
114. Lakoff, R.,
Language in Context, University of Michigan, 1970.
115. Lakoff, R.,
Politeness, Pragmatics, and Performatives, in **Proceedings of the Texas Conference on Performatives, Presuppositions, and Implicatures**, 1977.
116. Laucht, M. & Herrmann, T.,
Zur Direktheit von Direktiva, Universität Mannheim Technical Report no 1, 1978.
117. Lehnert, W.G.,
The Process of Question Answering, Hillsdale: Lawrence Erlbaum Associates, 1978.
118. Lehnert, W.G.,
Plot Units: A Narrative Summarization Strategy, in **Strategies for Natural Lan-**

- guage Processing**, Lehnert W.G. & Ringle M.H. (eds), Hillsdale: Lawrence Erlbaum Associates, 1982.
119. Levelt, W.J.M. & Schriefers, H.,
Stages of Lexical Access, in **Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics**, Kempen G. (ed), Boston: Kluwer Academic Publishers, 1987.
 120. Levinson, S.C.,
Pragmatics, Cambridge: Cambridge University Press, 1983.
 121. Levin, J.A. & Moore, J.A.,
Dialogue-Games: Metacommunication Structures for Natural Language Interaction, in *Cognitive Science* vol 1 no 4, 1977.
 122. Litman, D.J. & Allen, J.F.,
A Plan Recognition Model for Subdialogues in Conversations, University of Rochester Technical Report no 141, 1984.
 123. Loomis, R.S, Hull, H.R. & Robinson, M.L.,
The Art of Writing Prose, New York: Farrar & Rinehart, 1936.
 124. Lytinen, S. & Schank, R.C.,
Representation and Translation, Yale University Technical Report no 234, 1982.
 125. Mann, W.C.,
The Anatomy of a Systemic Choice, Information Sciences Institute Technical Report no RR-82-104 , 1982.
 126. Mann, W.C., 1983a,
An Overview of the Nigel Text Generation Grammar, Information Sciences Institute Technical Report no RR-83-113 , 1983.
 127. Mann, W.C., 1983b,
An Overview of the Penman Text Generation System, Information Sciences Institute Technical Report no RR-83-114 , 1983.

128. Mann, W.C.,
Discourse Structures for Text Generation, Information Sciences Institute Research Report RR-84-127, 1984.
129. Mann, W.C. & Thompson, S.A.,
Relational Propositions in Discourse, Information Sciences Institute Research Report RR-83-115, 1983.
130. Mann, W.C. & Thompson, S.A.,
Rhetorical Structure Theory: Description and Construction of Text Structures, in **Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics**, Kempen G. (ed), Boston: Kluwer Academic Publishers, 1987.
131. McCawley, J.D.,
Conversational Implicature and the Lexicon, in **Syntax and Semantics Vol 9: Pragmatics**, Cole P. (ed), New York: Academic Press, 1978.
132. McCoy, K.F.,
The Role of Perspective in Responding to Property Misconceptions, IJCAI Conference Proceedings, 1985.
133. McCoy, K.F.,
Contextual Effects on Responses to Misconceptions, in **Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics**, Kempen G. (ed), Boston: Kluwer Academic Publishers, 1987.
134. McDermott, D.V.,
Planning and Acting, in *Cognitive Science*, vol 2 no 2, 1978.
135. McDermott, D.V.,
Artificial Intelligence meets Natural Stupidity, in **Mind Design**, Haugeland J. (ed), Cambridge: MIT Press, 1981.

136. McDonald, D.D.,
Natural Language Production as a Process of Decision Making under Constraint,
 Ph.D. dissertation, Massachusetts Institute of Technology, 1980.
137. McDonald, D.D.,
Natural Language Generation as a Computational Problem: an Introduction, Univer-
 sity of Massachusetts (Amherst) Technical Report no 81-33, 1981.
138. McDonald, D.D. & Pustejovsky, J.D.,
Description-Directed Natural Language Generation, IJCAI Conference Proceedings,
 1985.
139. McGuire, R.,
Political Primaries and Words of Pain, unpublished manuscript, Yale University,
 1975.
140. McKeown, K.R.,
Generating Natural Language Text in Response to Questions about Database Queries,
 Ph.D. dissertation, University Of Pennsylvania, 1982.
141. McKeown, K.R.,
Recursion in Text and its Use in Language Generation, AAAI Conference Proceed-
 ings, 1983.
142. Michalski, R.,
The Theory and Methodology of Inductive Learning, in **Machine Learning**, Car-
 bonell J. & Mitchell T. (eds), San Francisco: Tioga Publishing Company, 1983.
143. Miller, D.P.,
Planning by Search Through Simulations, Ph.D. dissertation, Yale University, 1985.
144. Morris, C.W.,
Foundations of the Theory of Signs, in **International Encyclopedia of Unified
 Science**, vol 1, Neurath O., Carnap R. & Morris C.W. (eds), 1938.

145. Novak, H-J.,
A Relational Matching Strategy for Temporal Event Recognition, GWAI-84 Proceedings, Laubsch J. (ed), 1984.
146. Novak, H-J.,
Strategies for Generating Coherent Descriptions of Object Motions in Time-Varying Imagery, in **Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics**, Kempen G. (ed), Boston: Kluwer Academic Publishers, 1987.
147. Osgood, C.E.,
A Behavioristic Analysis of Perception and Language as Cognitive Phenomena, in **Contemporary Approaches to Cognition**, Bruner J. (ed), Cambridge: Harvard University Press, 1957.
148. Osgood, C.E. & Bock, J.K.,
Salience and Sentencing: Some Production Principles, in **Sentence Production: Developments in Research and Theory**, Rosenberg S. (ed), Hillsdale: Lawrence Erlbaum Associates, 1977.
149. Osgood, C.E., May, W.H. & Miron, M.,
Cross-Cultural Universals of Affective Meaning, Urbana: University of Illinois Press, 1975.
150. Paris, C.L.,
Descriptions Strategies for Naive and Expert Users, in Proceedings of the 23rd Annual Meeting of the Association for Computational Linguistics, Chicago, 1985.
151. Paris, C.L. & McKeown, K.R.,
Discourse Strategies for Descriptions of Complex Physical Objects, in **Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics**, Kempen G. (ed), Boston: Kluwer Academic Publishers, 1987.
152. Patten, T. & Ritchie, G.,
A Formal Model of Systemic Grammar, in **Natural Language Generation: Re-**

- cent **Advances in Artificial Intelligence, Psychology, and Linguistics**, Kempen G. (ed), Boston: Kluwer Academic Publishers, 1987.
153. Payne, L.V.,
The Lively Art of Writing, Chicago: Follett (Mentor Books), 1969.
 154. Piaget, J.,
The Language and Thought of the Child, New York: New American Library Inc. (Meridian Books), 1955.
 155. Reichman, R.,
Getting Computers to Talk like You and Me, Ph.D. dissertation, Harvard University, 1981.
 156. Reithinger, N.,
Antwortgenerierung für das Erlanger Spracherkennungssystem, Ph.D. dissertation, University of Erlangen-Nürnberg, 1984.
 157. Reithinger, N.,
Generating Referring Expressions and Pointing Gestures, in **Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics**, Kempen G. (ed), Boston: Kluwer Academic Publishers, 1987.
 158. Riesbeck, C.K. & Martin, C.E.,
Direct Memory Access Parsing, Yale University Technical Report no 354, 1985.
 159. Rich, E.A.,
User Modeling via Stereotypes, in *Cognitive Science*, vol 3 no 4, 1979.
 160. Rombauer, I.S. & Becker, M.R.,
The Joy of Cooking, New York: Bobbs-Merrill, 1975.
 161. Rosenberg, S.,
Semantic Constraints on Sentence Production: An Experimental Approach, in **Sentence Production: Developments in Research and Theory**, Rosenberg S. (ed), Hillsdale: Lawrence Erlbaum Associates, 1977.

162. Rösner D.,
The Automated News Agency SEMTEX — a Text Generator for German, in **Natural Language Generation: Recent Advances in Artificial Intelligence, Psychology, and Linguistics**, Kempen G. (ed), Boston: Kluwer Academic Publishers, 1987.
163. Sacerdoti, E.,
A Structure for Plans and Behavior, Amsterdam: North-Holland Publishing Company, 1977.
164. Sacks, H., Schegloff, E.A. & Jefferson, G.,
A Simplest Systematics for the Organization of Turn-Taking for Conversation, in *Language*, vol 50 no 4, 1974.
165. Schank, R.C.,
'Semantics' in Conceptual Analysis, in *Lingua* vol 30 no 2, Amsterdam: North-Holland Publishing Company, 1972.
166. Schank, R.C.,
Conceptual Information Processing, Amsterdam: North-Holland Publishing Company, 1975.
167. Schank, R.C.,
Rules and Topics in Conversation, in *Cognitive Science*, vol 1 no 1, 1977.
168. Schank, R.C.,
Interestingness: Controlling Inferences, in *Artificial Intelligence*, vol 12 no 3, 1979.
169. Schank, R.C.,
Dynamic Memory: A Theory of Reminding and Learning in Computers and People, Cambridge: Cambridge University Press, 1982.
170. Schank, R.C. & Abelson, R.P.,
Scripts, Plans, Goals and Understanding, Hillsdale: Lawrence Erlbaum Associates, 1977.

171. Schank, R.C. & Birnbaum, L.,
Memory, Meaning and Syntax, in **Talking Minds: Studies in the Cognitive Sciences**, Bever, Carroll & Miller (eds), Cambridge: MIT Press, 1983.
172. Schank, R.C., Collins, G.C., Davis, E., Johnson, P.N., Lytinen, S. & Reiser, B.J.,
What's the Point?, Yale University Technical Report no 205, 1981.
173. Schank, R.C. & Lehnert, W.,
The Conceptual Content of Conversation, Yale University Technical Report no 160, 1979.
174. Schegloff, E.A., Jefferson, G. & Sacks, H.,
The Preference for Self-Correction in the Organization and Repair of Conversation, in *Language*, vol 53 no 2, 1977.
175. Searle, J.R.,
Speech Acts, Cambridge: Cambridge University Press, 1969.
176. Searle, J.R.,
Indirect Speech Acts, in **Syntax and Semantics 3: Speech Acts**, Cole P. & Morgan J.L. (eds), New York: Academic Press, 1975.
177. Searle, J.R.,
Expression and Meaning, Cambridge: Cambridge University Press, 1979.
178. Sidner, C.L.,
Toward a Computational Theory of Definite Anaphora Comprehension in English, Massachusetts Institute of Technology Technical Report AI-TR-537, 1979.
179. Simmons, R.F. & Slocum, J.,
Generating English Discourse from Semantic Networks, in *Communications of the ACM* vol 15 no 10, 1972.
180. Stefik, M. & Bobrow, D.G.,
Object-Oriented Programming: Themes and Variations, in *AI Magazine* vol 6 no 4, 1986.

181. Stockwell, R.P., Schachter, P. & Partee, B.P.,
The Major Syntactic Structures of English, New York: Holt, Reinhart and Winston, 1972.
182. Straker, D.Y.,
Situational Variables in Language Use, University of Illinois (Urbana) Technical Report no 167, 1980.
183. Strunk, W. jr & White, E.B.,
The Elements of Style, New York: Macmillan, 1959.
184. Swartout, W.R.,
Producing Explanations and Justifications of Expert Consulting Programs, Massachusetts Institute of Technology Technical Report LCS-TR-251, 1981.
185. Sycara-Cyranski, K.,
Persuasive Argumentation in Resolution of Collective Bargaining Impasses, Cognitive Science Conference Proceedings, 1985.
186. Sycara-Cyranski, K.,
Arguments of Persuasion in Labour Mediation, IJCAI Conference Proceedings, 1985.
187. Van Dijk, T.,
Studies in the Pragmatics of Discourse, The Hague: Mouton, 1985.
188. Watzlawick, P., Beavin, J.H. & Jackson, D.D.,
Pragmatics of Human Communication, New York: Norton, 1967.
189. Weathers, W. & Winchester, O.,
The New Strategy of Style, New York: McGraw-Hill, 1978.
190. Wilensky, R.,
A Knowledge-Based Approach to Natural Language Processing: A Progress Report, IJCAI Conference Proceedings, 1981.

191. Wilensky, R.,
KODIAK - A Knowledge Representation Language Cognitive Science Conference Proceedings, 1984.
192. Willis, H.,
Structure, Style, and Usage (The Rhetoric of Exposition), (second edition),
New York: Holt, Rinehart & Winston, 1969.
193. Wish, Deutsch & Kaplan,
Perceived Dimensions of Interpersonal Relations, in *Journal of Personal and Social Psychology*, vol 33 no 4, 1976.
194. Wodehouse, P.G.,
Carry On, Jeeves, New York: Penguin, 1979.
195. Wong, H.K.T.,
Generating English Sentences from Semantic Structures, M.S. thesis, University of Toronto Technical Report no 84, 1975.
196. Woolf, B.P.,
Context Dependent Planning in a Machine Tutor, Ph.D. dissertation, University of Massachusetts (Amherst), 1984.
197. Woolf, B.P. & McDonald, D.D.,
Context-Dependent Transitions in Tutoring Discourse, AAAI Conference Proceedings, 1984.

END

9-87

DTIC