

An Introduction to Natural Language Generation

Robert Dale
 Microsoft Institute of Advanced Software Technology
 and
 School of Mathematics, Physics, Computing and Electronics
 Macquarie University
 Sydney
 Australia
 rdale@microsoft.com

-
1. An Overview of NLG
 2. Linguistic Realization
 3. Text Planning
 4. Generating Referring Expressions

Topic 3:

Text Planning

Strategy and Tactics in Generation

Strategic decisions:

- the purpose of an utterance
- the basic propositional material to be conveyed
- the relative importance of the elements in that material

Tactical decisions:

- the selection of lexical items
- the selection of syntactic constructions
- (in speech) the use of intonation and stress

Strategic Generation

Two foci of research:

- determining the content of a text
- determining the organization of a text

Overview

1. Davey's PROTEUS
2. McKeown's TEXT
3. Kukich's ANA
4. RST in NLG
5. Wahlster *et al*'s WIP

Davey's PROTEUS: An Example Output

The game began with your taking a corner, and I took the middle of an adjacent edge. You threatened me by taking the end of the one adjacent to the corner which you had just taken and to the one which I had just taken the middle of but I blocked that, and you took the square opposite the one I had just taken. I took the middle of the board and threatened you. If you had blocked my line, you would have threatened me, but you took the corner adjacent to the one which you took first and so I won by completing my line.


Davey's PROTEUS: An Example Output

Without the referring expressions:

- The game began with your taking ..., and I took
- You threatened me by taking ... but I blocked ..., and you took
- I took ... and threatened you.
- If you had blocked my ..., you would have threatened me, but you took ... and so I won by completing

Davey's PROTEUS

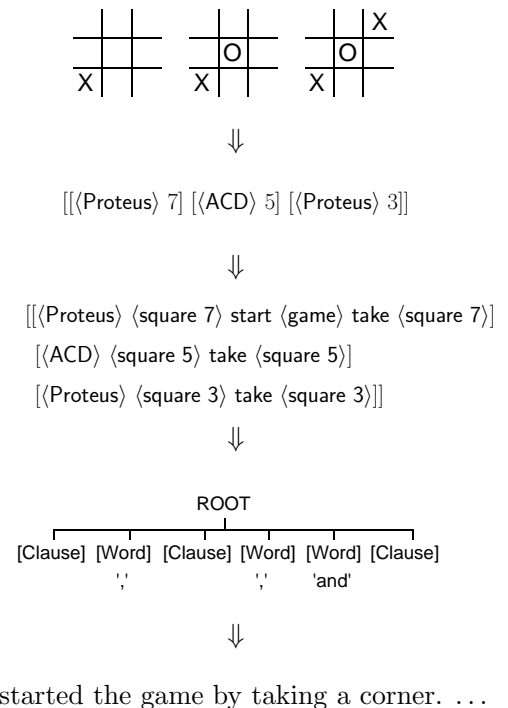
The Component Processes:



Davey's PROTEUS: How it Works

- the list of moves is annotated with evaluations and grouped into clusters
- the description of moves is decided upon, resulting in an intermediate representation which incorporates
 - (a) some assessment of the significance of the moves and
 - (b) some information required for building the sentence (the required punctuation and conjunctions)
- this representation is then used to produce the natural language output via a systemic grammar.

Davey's PROTEUS: An Example Input



Davey's PROTEUS

Move characterisation:

- each move is defensive, offensive or neutral
- each move is identified as a tactic (*blocking, forking, completing a win*)

Davey's PROTEUS

Annotations:

- priority-ordered list of heuristics that could result in this move
- tactically equivalent alternatives for each heuristic
- lines tactically involved in the move
- better move, if any

[From Ritchie's 1984 reconstruction.]

Davey's PROTEUS

The game began with your taking a corner, and I took the middle of an adjacent edge. If you had taken the corner opposite the one which you had just taken, you would have threatened me, but you took the one adjacent to the square which I had just taken.

[[⟨ACD⟩ 1] [⟨Proteus⟩ 2] [⟨ACD⟩ 3]]

Annotations on third move:

Square	3
Heuristics	take [9 8 7 6 5 4]
Better Move	Square 9 (1 5 9)
	Heuristics threaten [7 (1 4 7) 5 (1 5 9) 4 (1 4 7)]

Davey's PROTEUS

Sentence scope heuristics:

1. If there is one move left, express as a single sentence.
2. If there are two moves left, express as a single sentence.
3. If a move is a mistake, start a new sentence to describe it.
4. If a move is a combined attack and defence give it a sentence to itself.
5. If this move is an attack and the next move thwarts this attack put these two moves into a sentence.
6. Put the next three moves into a sentence.

Davey's PROTEUS

Coherence through sequentiality:

- Move-1, and Move-2
You started the game by taking a corner, and I took the opposite one.
- Move-1, Move-2, and Move-3
The game began with my taking a corner, you took an adjacent one, and I took the middle of the same edge.

Davey's PROTEUS

Coherence through contrast:

- Move-1 but Move-2
I threatened you by taking the middle of the board but you blocked my line.
- Move-1 but Move-2. However, Move-3.
I threatened you by taking the middle of the board, but you blocked my line and threatened me. However, I blocked your edge by taking the middle of it.

Davey's PROTEUS

Avoiding surprise:

- Although Move-1, Move-2.
Although you blocked one of my edges, I won by completing the other.

Mentioning mistakes:

- Hypothetical but Move-1.
You could have forked me but you took the square opposite the one I had just taken.
- Conditional-Subordinate but Move-1.
You could have forked me but you took the square opposite the one I had just taken and so I won by completing my edge.

Davey's PROTEUS

Subordination:

- You started the game by taking a corner.
- I threatened you by blocking your line.

Logical Order:

- I blocked your line and threatened you.
- * I threatened you and blocked your line.

Davey's PROTEUS

Summary:

- very impressive for its time: fluency still superior to many current NLG systems
- the beginnings of a computational account of discourse coherence
- techniques used inseparable from domain knowledge

Overview

1. Davey's PROTEUS
2. McKeown's TEXT
3. Kukich's ANA
4. RST in NLG
5. Wahlster *et al*'s WIP

McKeown's TEXT

Observations:

- people generally follow standard patterns of organization when producing texts.

Hypothesis:

- people have preconceived ideas about what resources to use to achieve particular goals;
- people have preconceived ideas about how these resources can be integrated to form a text.

So:

- if these patterns are captured formally, they can guide a generation system in its decisions about what to say next.

McKeown's TEXT

Developed to provide paragraph-length responses to meta-level questions about the structure of an underlying database.

Could generate responses to three classes of questions—three COMMUNICATIVE GOALS:

- requests to define a concept
- requests to compare two objects
- requests to describe available information

The Basic Ideas

- Each utterance in a discourse plays some role in that discourse—characterised by a RHETORICAL PREDICATE.
- Examples: making an analogy, describing sub-parts or sub-types, providing detail about something.
- These rhetorical predicates can be combined into standard patterns of discourse represented by SCHEMATA.

Rhetorical Predicates

Attributive: *Mary has a pink coat.*

Equivalent: *Wines described as 'great' are fine wines from an especially good village.*

Specification: *[Mary is quite heavy.] She weighs 200 pounds.*

Constituency: *[This is an octopus.] There is his eye, these are his legs, and he has these suction cups.*

Evidence: *[The audience recognized the difference.] They started laughing right from the very first frames of that film.*

...

Adversative: *It was a case of sink or swim.*

Inference: *So people form a low self-image of themselves.*

McKeown's TEXT

McKeown

- analysed naturally occurring texts for the three DISCOURSE PURPOSES
- identified four DISCOURSE STRATEGIES that together accounted for the structure of paragraphs

Purpose	Strategy
define	IDENTIFICATION CONSTITUENCY
compare	COMPARE AND CONTRAST
describe	ATTRIBUTIVE CONSTITUENCY

The Attributive Schema

Attributive

{Amplification; Restriction}

Particular illustration*

{Representative}

{Question; Problem

Answer} /

{Comparison; Contrast

Adversative}

Amplification/Explanation/Inference/Comparison

The Attributive Schema

Attributive: This book, being about work, is, by its very nature, about violence—

Amplification: to the spirit as well as to the body.

Particular Illustration: It is about ulcers as well as accidents, about shouting matches as well as fistfights, about nervous breakdowns as well as kicking the dog around.

Representative: It is, above all (or beneath all), about daily humiliations.

Amplification; Explanation: To survive the day is triumph enough for the walking wounded among the great many of us.

The Identification Schema

Identification (class and attribute function)

{Analogy/Constituency/Attributive/

Renaming/Amplification}*
Particular illustration/Evidence +

{Amplification/Analogy/Attributive}

{Particular illustration/Evidence}

The Identification Schema

Eltville (Germany):

Identification: An important wine village of the Rheingau region.

Attributive: The vineyards make wines that are emphatically of the Rheingau style

...

Amplification: ... with a considerable weight for a white wine.

Particular Illustration: Taubenberg, Sonnenberg and Langenstuck are among the vineyards of note.

The Constituency Schema

Constituency

Cause-effect*/Attributive*/

{Depth-identification/Depth-attributive

{Particular-illustration/Evidence}

{Comparison/Analogy}} +

{Amplification/Explanation/Attributive/Analogy}

The Constituency Schema

Steam and electric torpedoes:

Constituency: Modern torpedoes are of two general types.

Depth-identification; (Depth-attributive): Steam-propelled models have speeds of 27 to 45 knots and ranges of 4000 to 25000 yards.

Comparison: The electric powered models are similar ...

Depth-identification; (Depth-attributive): ... but do not leave the telltale wake created by the exhaust of a steam torpedo.

The Compare and Contrast Schema

Positing/Attributive ($\neg A$)

{Attributive (A)/

Particular-illustration/Evidence (A)/

Amplification (A)/

Inference (A)/Explanation (A)}+

{Comparison ($A \wedge \neg A$)/Explanation ($A \wedge \neg A$)/

Generalization ($A \wedge \neg A$)/Inference ($A \wedge \neg A$)}

The Compare and Contrast Schema

Positing ($\neg A$): Movies set up these glamourized occupations.

Attributive (A): When people find they are waitresses, they feel degraded.


Evidence (A): No kid says I want to be a waiter, I want to run a cleaning establishment.

Comparison, Explanation ($A \wedge \neg A$): There is a tendency in movies to degrade people if they don't have white-collar professions.

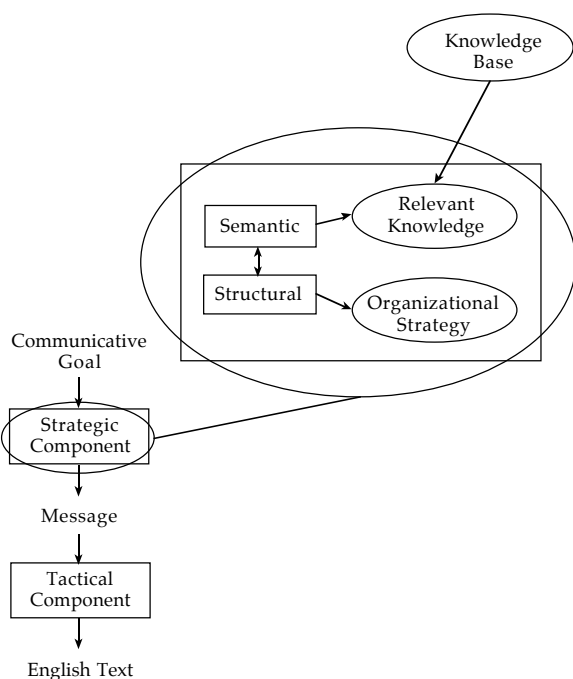
Inference ($A \wedge \neg A$): So, people form a low self-image of themselves ...

Comparison, Explanation ($A \wedge \neg A$): ... because their lives can never match the way Americans live—on screen.

McKeown's TEXT: Overall Structure



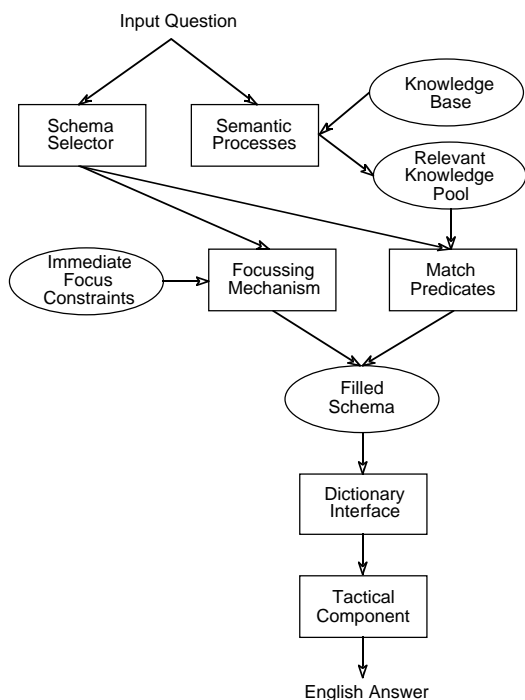
McKeown's TEXT: Overall Structure



McKeown's TEXT: How it Works

- Choose set of schemata on the basis of question.
- Select relevant knowledge pool on the basis of question.
- Select one schema on the basis of the knowledge pool.
- Fill the schema via the semantics of the rhetorical predicates.
- Where there are alternatives in the schema, use focus of attention to select the best alternative.
- Realize utterances via a functional grammar.

McKeown's TEXT



Choosing a Set of Schemata

Purpose Strategy

define	IDENTIFICATION CONSTITUENCY
compare	COMPARE AND CONTRAST
describe	ATTRIBUTIVE CONSTITUENCY

Determining the Relevant Knowledge

Requests for information or definitions:

Section off area around the questioned object; preserve all links; include siblings and descendents of questioned object.

Comparisons: If the two entities are similar, need to provide detail; if they are very different, need discussion of generic class.

Determining the Relevant Knowledge

Comparisons:

1. What is the difference between a part-time and a full-time student?
2. What is the difference between a raven and a writing desk?
3. A part-time student takes 2 or 3 courses per semester while a full-time student takes 3 or 4.
4. A writing desk has 4 legs while a raven has only 2.

Selecting One Schema

Requests for information or definitions:

- if the knowledge pool contains a rich description of the object's sub-classes use the **Constituency** schema
- if there is more information about the object itself use the **Identification** schema for definitions and the **Attributive** schema for information questions.

Filling the Schema

- Each predicate has semantics which define what that predicate means in the database domain.
- Each predicate is a type of **inform** speech act.
- A proposition is an instantiation of a predicate, assembled from the relevant knowledge pool.

The Attributive Predicate

Given Argument: entity

Type: DB-attributes

Subtype: attributes-only

```
(attributive db <entity> <naming-attr>
  <topic-attr> <duplicate-attr> <db-attribs>)
```

Example:

```
(attributive db SHIP (name NAME)
  (topics DIMENSIONS)
  (duplicates
    (FUEL (FUEL-TYPE FUEL-CAPACITY))
  (attrs MAXIMUM-SPEED))
```

Using Focus

- Where there are alternatives in the schema, focus of attention is used to select the best alternative.
- McKeown's ordering on Sidner's focussing constraints causes the system to shift attention to an item just introduced if possible, otherwise to maintain focus, and in all other cases to return to an earlier focus of attention.
- An indication of current focus can also be passed to the tactical component to assist in realization decisions.

McKeown's TEXT

identification
 constituency
 attributive
 attributive
 evidence
 evidence
 attributive

1. A guided missile is a projectile that is self-propelled.
2. There are two types of guided projectiles in the ONR database: torpedoes and missiles.
3. The missile has a target location in the air or on the earth's surface.
4. The torpedo has an underwater target location.
5. The missile's target location is indicated by the DB attribute DESCRIPTION and the missile's flight capabilities are provided by the DB attribute ALTITUDE.
6. The torpedo's underwater capabilities are provided by the DB attributes under DEPTH (for example, MAXIMUM OPERATING DEPTH).
7. The guided projectile has DB attributes TIME TO TARGET & UNITS, HORZ RANGE & UNITS and NAME.

McKeown's TEXT

Summary:

A flexible means of generating a restricted range of texts:

- since a single discourse purpose may be realized by more than one strategy, the structure of a generated text will differ depending upon the strategy chosen
- since any single strategy has a number of alternatives, the resulting text will be different depending upon which alternatives are taken

McKeown's TEXT

Limitations:

- doesn't obviously scale-up to the full variety of general purpose text generation
- the system does not take into account any constraints deriving from information about the user
- the semantics of the rhetorical predicates hold only for the database domain

Overview

1. Davey's PROTEUS
2. McKeown's TEXT
3. Kukich's ANA
4. RST in NLG
5. Wahlster *et al*'s WIP

Ana: Generating Text from Databases

Kukich 1985:

- The problem: information overload
- One solution: need to convert machine-readable data into human-readable information
- One method: knowledge-based report generation to produce summaries of numerical data

Some Input Data

06/24	IND	814.69	821.63	805.56	810.41	DN	2.76
06/24	TRN	317.97	321.57	313.85	317.00	UP	0.30
06/24	UTL	106.13	106.83	105.61	106.83	UP	0.70
06/24	CLOSINGSTOCK		30	INDUS	810.41		
06/24	CLOSINGSTOCK		20	TRANSP	317.00		
06/24	CLOSINGSTOCK		15	UTILS	106.83		
06/24	CLOSINGSTOCK		65	STOCKS	314.90		
06/24	330PM		30	INDUS	812.02		
06/24	330PM		20	TRANSP	317.60		
06/24	330PM		15	UTILS	106.52		
06/24	330PM		65	STOCKS	315.28		
06/24	3PM		30	INDUS	810.12		
06/24	3PM		20	TRANSP	316.47		
06/24	3PM		15	UTILS	106.70		
06/24	3PM		65	STOCKS	314.63		
06/24	230PM		30	INDUS	810.50		
06/24	230PM		20	TRANSP	316.77		
...							
06/24	1130AM		15	UTILS	106.57		
06/24	1130AM		65	STOCKS	315.81		
06/24	11AM		30	INDUS	813.55		
06/24	11AM		20	TRANSP	317.60		
06/24	11AM		15	UTILS	106.22		
06/24	11AM		65	STOCKS	315.46		
06/24	1030AM		30	INDUS	814.12		
06/24	1030AM		20	TRANSP	318.05		
06/24	1030AM		15	UTILS	106.39		
06/24	1030AM		65	STOCKS	315.79		
06/24	ISSUETRADED			1849			
06/24	ADVANCED			772			
06/24	DECLINED			660			
06/24	UNCHANGED			417			
06/24	NATION			65223320	72820880		
06/24	NYSEVOL			55860000	62710000		
06/23	IND	801.75	815.35	795.76	813.17	UP	13.51
06/23	TRN	308.98	317.37	306.65	316.70	UP	8.92
06/23	UTL	106.13	106.83	105.61	106.13	DN	0.17

A Human-Generated Summary

Stocks retreated from a broad advance yesterday and closed mixed. Trading was heavy.

Retail, utility and forest-products issues gained ground, but chemical, auto, drug, oil, and other energy-related stocks were numerous among the losers.

The Dow Jones average of thirty industrials surrendered a 7.23 gain at 2pm EST and closed at 825.82, down 2.57. The transportation and utility measures edged higher ...

[Pittsburgh Post-Gazette, March 3, 1982]

Deriving Information from Data

1. Damon was up $1/8$ at $7 \frac{1}{2}$.
Dan River was down $3/8$ to $13 \frac{7}{8}$.
Data General was down $2 \frac{1}{8}$ to $50 \frac{1}{2}$.
Dayco Corp was up $1/8$ at 11.
Industrials were down 1.90 at 868.91.
Utilities were down 0.32 at 108.61.
2. Ford managed a $1/8$ gain to $18 \frac{1}{2}$, but General Motors fell $1/4$ to 38, American Motors fell $1/8$ to $2 \frac{5}{8}$, and Chrysler slipped $1/8$ to $4 \frac{1}{4}$.
3. Even brisk trading in AT&T and IBM stock was unable to stimulate the Dow Jones average of 30 industrials.

Expressing Information Fluently

1. The Dow Jones utilities average hit a new low. The Dow Jones transportation average hit a new low. The American Stock Exchange's market value index hit a new low. The NYSE's composite index hit a new low. The NYSE's index is a composite of more than 1500 common stocks. The NYSE's composite index fell 1.00 to 64.14.
2. But several other indicators hit new lows, among them the NYSE's composite index of more than 1500 common stocks, which fell 1.00 to 64.14.

Expressing Information Fluently

Decisions:

- which words and syntactic forms should be used to express an idea
- how to combine words or phrases into well-formed sentences
- how to combine sentences into longer sentences
- when to eliminate redundancy by replacing nouns with pronouns or introducing ellipsis
- how to maintain focus and coherence within the structure of the text.

Key Ideas in ANA

- sublanguage
- the phrasal lexicon

Sublanguage

The written material in a specialised technical area displays linguistic regularities over and above those that can be stated for the language as a whole, giving rise to the notion of sublanguage. Sublanguage texts can be characterised by the occurrence of word classes unique to the discipline, and by the syntactic combinations in which those classes occur.

[Sager 1981:199]

These regularities define semantic and linguistic knowledge requirements for natural language processing systems.

Sublanguage and Knowledge Representation

...the sublanguage grammar is more than just a linguistic characterisation of the texts. The lexical classes and the hierarchical relations between the classes usually reflect the accepted taxonomy which the specialised field of knowledge imposes on the objects of its limited domain of discourse. And the combinations of lexical classes which are permissible in the sentences of the specialised texts reflect conceivable relations between these objects. Thus the sublanguage grammar can be said to incorporate certain aspects of a knowledge representation for the subfield.

[Kittredge and Lehrberger 1983:8]

Core-Market *vs* Extra-Market Content

Kukich's analysis of 24 stock reports:

- of the 354 sentences, 72% expressed ideas inherent in the stock market data: messages about indicator changes, trading levels, historical comparisons
- 28% referred to events external to the stock market: news reports of decrease-oil production, union contract negotiation, federal economic policies

Example:

Eastman Kodak advanced 2 3/4 to 85 3/4; it announced development of the world's fastest color film for delivery in 1983.

The Phrasal Lexicon

I suggest that utterances are formed by the repetition, modification, and concatenation 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; productive processes have the secondary role of adapting old phrases to the new situation.

...

We start with the information we wish to convey and the attitudes toward that information that we wish to express or evoke, and we haul out of our phrasal lexicon some patterns that can provide the major elements of this expression. Then the problem is to stitch these phrases together into something roughly grammatical, to fill in the blanks with the particulars of the case at hand, to modify the phrases if need be, and if all else fails to generate phrases from scratch to smooth over transitions or fill in any remaining conceptual holes.

[Becker 1975:70]

Actual Output from ANA

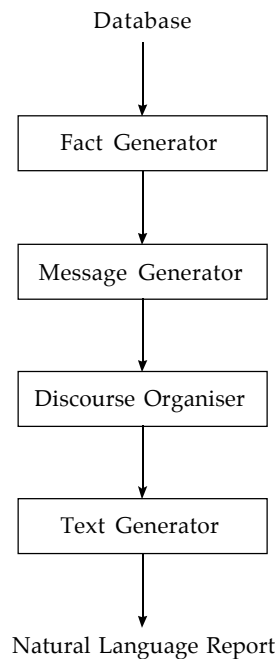
Thursday June 24, 1982

wall street's securities markets meandered upward through most of the morning, before being pushed downhill late in the day yesterday. the stock market closed out the day with a small loss and turned in a mixed showing in moderate trading.

the Dow Jones average of 30 industrials declined slightly, finishing the day at 810.41, off 2.76 points. the transportation and utility indicators edged higher.

volume on the big board was 55860000 shares compared with 62710000 shares on Wednesday. advances were ahead by about 8 to 7 at the final bell.

Ana's Architecture



The Fact Generator

- Example: ... *meandered upward through most of the morning*
- Fact Generator uses the half-hour interval data to determine degree and direction of change for the half-hour and cumulative degree and direction of change for the day

The Fact Generator

A Half-Hour Statistics Fact:

```

FactName: HRStat
IndicatorName: DowJonesIndustrial
IndicatorType: Composite
Date: 04-21
Hour: 1130am
NHour: 1130
CurrentLevel: 1192.82
DirectionOfChange: Down
DegreeOfChange: 2.03
CumulativeDirection: Up
CumulativeDegree: 1.35
HighOrLowMark: nil
  
```

The Fact Generator

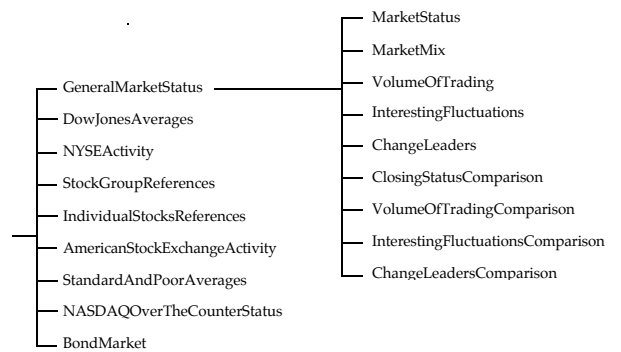
HRStats are just one of four kinds of facts. Others:

- CIStats: closing statistics for composite indicators
- VolStats: trading volume statistics
- AVDDDECs: advance-decline facts—the number of issues advancing, declining, unchanged and traded, plus the computed ratio of advancing to declining stocks

The Message Generator

- builds interesting messages from facts
- analysis of stock market reports suggests messages can be classified in a taxonomy
- a set of inference rules recognize patterns in facts and instantiate semantic attributes of messages accordingly

The Message Hierarchy



The Message Generator

Describing Market Status:

- The stock market closed out the day with a small loss.
- The stock market closed slightly lower yesterday.
- The stock market staged a moderate advance yesterday.
- The stock market trembled yesterday.
- Stock prices crept upward yesterday.
- The stock market slipped into a broad decline yesterday.

The Message Generator

Semantic attributes of market status messages:

- market direction
- degree of change
- duration of change
- scope of change
- time of day

The market was swept into a broad and steep decline.

direction: down
 degree: great
 scope: broad
 time: close

The Message Generator

A Dow Status Message:

The Dow Jones average of 30 industrials declined slightly, finishing the day at 810.41, off 2.76 points.

The Underlying Message Structure:

```

ReportDate: 06/24
Topic: Dow
SubTopic: DowAverageStatus
SubSubTopic: nil
SubjectClass: Dow
Direction: Down
Degree: Small
Time: Close
VariableLevel: 810.41
VariableDegree: 2.76
  
```

The Message Generator

Rule 1:

if there is an active goal
 to instantiate the dow status message
then
 create a partially instantiated dow message
 create four new goals:
 instantiate the objective attributes
 instantiate the direction attribute
 instantiate the degree attribute
 write the message to the output file

The Message Generator

Rule 2:

if there is an active goal
 to instantiating objective variables
and there is a dow status message
and there is a fact indicating that
 the closing status of the dow
 on Date was Level and Degree
then modify the dow status message
 to instantiate the Date, Level and Degree

The Message Generator

Rule 3:

if there is an active goal
 to instantiate the direction attribute
and there is a message whose status is nil,
 whose topic is either general market or dow
and whose subtopic is either market status
 or dow status
and there is a dow closing status fact
 whose cumulative direction is Direction
then modify the message
 to instantiate the Direction

The Message Generator

Rule 4:

if there is an active goal
 to instantiate the degree-of-change attribute
and there is either a market status
 or dow status message whose status is nil
and there is a closing status fact
 for the dji average whose cumulative
 degree-of-change is greater than 10
then set the degree-of-change attribute at great

The Discourse Organiser

Functions of the Discourse Organiser:

- take messages produced by the message generator
- group into paragraphs by topic
- order within paragraphs

The Discourse Organiser

The Default Discourse Structure:

1. General Market Status
 - (a) Interesting Fluctuations
 - (b) Change Leaders
 - (c) Market Status
 - (d) Market Mix
 - (e) Volume of Trading
2. Dow Jones Averages
3. NYSE Activity
4. Stock Group References
5. Individual Stocks References
6. American Stock Exchange Activity
7. Standard and Poor Averages
8. NASDAQ Over The Counter Status
9. Bond Market

Discourse Organisation Productions

Some pragmatic rules which may revise or extend the default structure:

- for refining the default structure when default knowledge is insufficient;
- for combining or eliminating messages when similarities or redundancies are detected;
- for revising the default structure under exceptional circumstances.

The Text Generator

- focus on a message
- select a phrase that matches the semantic attributes of the message from a phrasal dictionary
- select a valid and rhetorically appropriate syntactic form
- combine with the foregoing clauses to produce fluent text

Messages are processed linearly, with one message lookahead.

Phrasal Lexical Entries

A message:

```
Priority: 1
ReportDate: 06/24
Topic: GeneralMarket
SubTopic: MktPt
SubSubTopic: MornUpSmLn
SubjectClass: Market
Direction: Up
Degree: Small
Duration: Long
Time: Early
```

Phrasal Lexical Entries

A phrasal lexicon entry:

```
PhraseType: Predicate
Topic: GeneralMarket
SubTopic: MktPt
Direction: Up
Degree: Small
Duration: Long
Time: Early
SubjType: Name
SubjectClass: Market
PredFiniteSingular: crept
PredFinitePlural: crept
PredParticipial: creeping
PredInfinitive: |to creep|
PredRemainder: |upward early in the morning|
Len: 12
Rand: 5
```

Kinds of Knowledge

- lexical, semantic, syntactic and rhetorical knowledge
- Len and Rand are ‘rhetorical’ attributes.
- Len = length in syllables of the phrase
- There is a rhetorical constraint that says ‘generate twolong sentences followed by one short one’, where long is 30 syllables
- Rand = a preference weighting attribute—higher values are chosen when a random choice is made

Another Lexical Entry

A phrasal lexicon entry:

```
PhraseType: Subject
Topic: GeneralMarket
SubjectClass: Market
SubjType: Name
SubjTerm: |the stock market|
SubjNumber: singular
Rand: 9
SubjHyponymLevel: 2
Usage: 0
```

Rhetorical Information

- Usage is used to avoid undesirable repetition
- Hyponym values are used by subject selector rules to ensure the selection of more familiar terms when subjects are repeatedly required from the same subject class within the same paragraph:

the Dow Jones average of 30 industrials ... the industrials average ... the Dow

Summary

- ANA generates natural language summaries from an underlying database
- deriving interesting facts from data is key
- empirical phenomenon of sublanguage justifies limited domain and drives knowledge representation
- phrasal lexicon avoids need for complex grammatical processing at the cost of flexibility
- discourse structures adopted are relatively rigid

Overview

1. Davey's PROTEUS
2. McKeown's TEXT
3. Kukich's ANA
4. RST in NLG
5. Wahlster *et al*'s WIP

A Reassessment

- schema-based approaches too inflexible for general-purpose text generation
- we need some better notion of coherence in text in order to generate coherent texts
- a better approach: the dynamic paragraph planning paradigm:
 - don't use fixed schemas
 - try to build coherent text from atomic units and rhetorical relations

Rhetorical Structure Theory

Questions to be answered:

- what are the discernible parts of an organized text?
- how can these parts be arranged?
- how can the parts be connected together to form a whole text?

Rhetorical Structure Theory

The basic idea:

- embodies a set of rhetorical schemata for a wide variety of purposes
- an approach to what it means for a text to be coherent

Aim of the theory:

- the development of a comprehensive theory of text organization that can eventually be used as part of a text generation system

Coherence

Where does coherence come from?

- (1) I love to collect classic automobiles.
- (2) a. My favourite car is my 1899 Duryea.
b. My favourite car is my 1987 Ford Escort.

Observations:

- The two texts differ in coherence.
- The difference is due to the 1899 Duryea being a classic car but not the 1987 Ford Escort.
- There is no explicit statement that either car is a classic car.

The Basic Ideas behind RST

Rhetorical Structure Theory:

- a descriptive theory that characterizes the structure of texts in terms of functional relations that hold between parts of the text;
- identifies the organizational resources available to the writer;
- views texts as hierarchical in structure.

Scope:


- The theory is intended to be insensitive to text size.
- Current theory is restricted to monologue.

Relational Propositions


(1) I love to collect classic automobiles.

My favourite car is my 1899 Duryea.

elaboration(love-to-collect, favourite-car)



Relation Definitions




A RELATION holds between two portions of a text, called the NUCLEUS and the SATELLITE (N and S).

A Relation Definition: ELABORATION

<i>relation name:</i>	ELABORATION
<i>constraints on N:</i>	none
<i>constraints on S:</i>	none
<i>constraints on the N + S combination:</i>	S presents additional detail about the situation or some element of subject matter which is presented in N or inferentially accessible in N in one or more of the ways listed below. In the list, if N presents the first member of any pair, then S includes the second:
	<ol style="list-style-type: none"> 1. set:member 2. abstract:instance 3. whole:part 4. process:step 5. object:attribute 6. generalisation:specific
<i>the effect:</i>	R recognizes the situation presented in S as providing additional detail for N. R identifies the element of subject matter for which detail is provided.
<i>the locus of the effect:</i>	N and S.

A Relation Definition: JUSTIFY



relation name JUSTIFY
constraints on N none
constraints on S none
constraints on the N + S combination:
 R's comprehending S increases R's
 readiness to accept W's right to
 present N
the effect R's readiness to accept W's right to
 present N is increased
the locus of the effect N

Another Example: Motivation

D3-1 You should come to the Los Angeles Chamber Ballet (the ballet company I'm dancing with) next week.

D3-2 Tickets are \$7.50 except for the opening night.


D3-3 The show is made up of new choreography and should be very interesting.

D3-4 I'm in three pieces.

enablement(D3-1, D3-2)

motivation(D3-1, D3-3)

motivation(D3-1, D3-4)



The MOTIVATION Relation Definition


<i>relation name:</i>	MOTIVATION
<i>constraints on N:</i>	presents an action on which R is the actor, unrealized with respect to the context of N.
<i>constraints on S:</i>	none
<i>constraints on the N + S combination:</i>	Comprehending S increases R's desire to perform action presented in N.
<i>the effect:</i>	R's desire to perform the action presented in N is increased.
<i>the locus of the effect:</i>	N

Another Example: Evidence

D4-1 The program as published for calendar year 1980 really works.

D4-2 In only a few minutes, I entered all the figures from my 1980 tax return and got a result which agreed with my hand calculations to the penny.

evidence(D4-1, D4-2)



The EVIDENCE Relation Definition

<i>relation name:</i>	EVIDENCE
<i>constraints on N:</i>	R might not believe N to a degree satisfactory to W
<i>constraints on S:</i>	R believes S or will find it credible.
<i>constraints on the N + S combination:</i>	R's comprehending S increases R's belief of N.
<i>the effect:</i>	R's belief of N is increased.
<i>the locus of the effect:</i>	N

Rhetorical Structure Theory

Nuclei and satellites:

- in most cases, the function of a satellite is to support the nucleus in some way
- there is a RELATION between each satellite and the nucleus indicating how the satellite provides support
- the schemata are unordered: i.e., the satellite and the nucleus can appear in any order in the text
- schemata are also recursive: a text span serving as the satellite of one schema may be decomposed into a nucleus and satellite of its own using another schema

Schema Application Conventions

- one schema is instantiated to describe the entire text
- schemas are instantiated to describe the text spans produced in instantiating other schemas
- the schemas do not constrain the order of nucleus or satellites in the text span in which the schema is instantiated; but nucleus and satellites tend to occur in certain orders
- a relation which is part of a schema may be instantiated indefinitely many times in the instantiation of that schema; but relations tend to be used just once

The Set of Relations

- Circumstance
- Solutionhood
- Elaboration
- Background
- Enablement and Motivation
- The "Cause" Cluster: Cause, Result and Purpose
 - Volitional Cause
 - Non-volitional Cause
 - Volitional Result
 - Non-volitional Result
 - Purpose
- Antithesis and Concession
- Condition and Otherwise
- Interpretation and Evaluation
- Restatement and Summary
- The Multi-Nuclear Relations
 - Sequence
 - Contrast
 - Joint

Analyses Performed with RST

100s of texts have been analysed using RST.

administrative memos, magazine articles, advertisements, personal letters, political essays, scientific abstracts, ...


Outcomes of RST Analysis:

- virtually every text has an RST analysis
- certain text types characteristically don't have RST analyses: for example, laws, contracts, and various kinds of "language-as-art" such as poetry
- in our culture, texts that have RST analyses predominate; it is thus typical for texts to be hierarchically structured and functionally organized.

A Larger Example Analysis

- 1 What if you're having to clean floppy drive heads too often?
- 2 Ask for SYNCOM diskettes, with burnished Ectype coating and dust-absorbing jacket liners.
- 3 As your floppy drive writes or reads,
- 4 a SYNCOM diskette is working four ways
- 5 to keep loose particles and dust from causing soft errors, dropouts.
- 6 Cleansing agents on the burnished surface of the Ectype coating actually remove build-up from the head,
- 7 while lubricating it at the same time.
- 8 A carbon additive drains away static electricity
- 9 before it can attract dust or lint.
- 10 Strong binders hold the signal-carrying oxides tightly within the coating.
- 11a And the non-woven jacket liner,
- 12 more than just wiping the surface,
- 11b provides thousands of tiny pockets to keep what it collects.
- 13 To see which SYNCOM diskette will replace the ones you're using now,
- 14 send for our free "Flexi-Finder" selection guide and the name of the supplier nearest you.
- 15 Syncom, Box 130, Mitchell, SD 57301. 800-843-9862; 605-996-8200.

A Larger Example Analysis



What is Nuclearity?

- Commonalities observed in the asymmetries between pairs of text spans: often, one member of the pair is
 - incomprehensible independent of the other, a non-sequitur, but not vice versa;
 - more suitable for substitution than the other;
 - more essential to the writer's purpose than the other.
- Removing the most-nuclear element makes the text unclear or even incoherent.
- Delete all the leaf-node satellites and you get something like a synopsis; delete all the nuclei and you get incoherence.

Deleting Satellites

- 2 Ask for SYNCOM diskettes, with burnished Ectype coating and dust-absorbing jacket liners.
- 4 a SYNCOM diskette is working four ways
- 6 Cleansing agents on the burnished surface of the Ectype coating actually remove build-up from the head,
- 8 A carbon additive drains away static electricity
- 10 Strong binders hold the signal-carrying oxides tightly within the coating.
- 11 And the non-woven jacket liner, . . . provides thousands of tiny pockets to keep what it collects.
- 14 send for our free "Flexi-Finder" selection guide and the name of the supplier nearest you.

Deleting Nuclei

- 1 What if you're having to clean floppy drive heads too often?
- 3 As your floppy drive writes or reads,
- 5 to keep loose particles and dust from causing soft errors, dropouts.
- 7 while lubricating it at the same time.
- 9 before it can attract dust or lint.
- 12 more than just wiping the surface,
- 13 To see which SYNCOM diskette will replace the ones you're using now

Rhetorical Structure Theory: A Summary

Elements of the theory:

- RST consists of a set of schemata
- each schema indicates how a unit of text structure can be decomposed into smaller units
- each such unit is a TEXT SPAN
- a schema consists of a NUCLEUS and zero or more SATELLITES

What We Want for Computational Implementation

generativity: capable of use in text construction as well as text description

functionality: informative in terms of how text achieves its effects for the writer

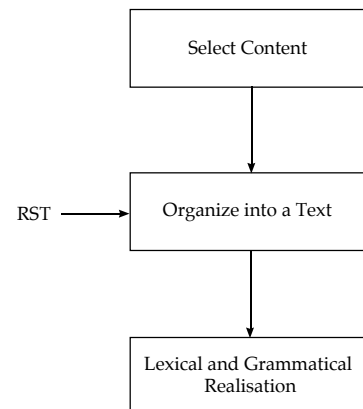
scale insensitivity: applicable to every size of text, and capable of describing all the various sized units of text organisation that occur

definiteness susceptible to formalization and programming

How We Might Use RST in Generation

- Before building RST structure
 - general decisions about what to accomplish, what knowledge to use; this yields a body of material to convey
- RST structure building, including
 - organizing the given body of material
 - supplementing it as needed, with evidence, concessives, circumstantials, antithesis, contrast and other supporting material
- After building RST structure
 - theme control, sentence scope, conjunction uses, lexical choice, formulaic text (e.g., “Sincerely yours”), grammatical realization

Where RST Might Fit in the Generation Process



Hovy's RST Text Structurer

What the text structurer does:

- accepts inputs from the domain of discourse;
- rewrites into a common form called INPUT UNITS;
- assembles input units into a tree that expresses paragraph structure; and
- dispatches leaves of tree to be generated by a linguistic realization component.

Assumption: the information to be conveyed has already been selected.

Hovy's RST Text Structurer

Typical set of input elements:

```

((SHIP.EMPLOYMENT A105)
 (SHIP.R A105 KNOX)
 (SHIP.COURSE.R A105 195)
 (CURRENT.POSITION.R A105 P102)
 (POSITION P102)
 (LONGITUDE.R P102 79)
 (LATITUDE.R P102 18)
 (READINESS.LEVEL.R A105 C4)
 (NEXT.MAJOR.EMPLOYMENT.R A105 E107)
 (CURRENT.MAJOR.EMPLOYMENT.R A105 E105)
 (ENROUTE E105)
 (EBEG.R E105 870420)
 (EEND.R E105 870424)
 (DESTINATION.R E105 SASEBO)
 (LOAD E107)
 (EBEG.R E107 870425)
 (EEND.R E107 870428)
  
```

Hovy's RST Text Structurer

The resulting input units:

```
((ENROUTE E105)
(SHIP.R E105 KNOX)
(DESTINATION.R E105 SASEBO)
(HEADING.R E105 HEADING11416)
(READINESS.R E105 READINESS11408)
(NEXT-ACTION.R E105 ARRIVE11400))

((LOAD E107)
(SHIP.R E107 KNOX)
(EBEG.R E107 870425)
(EEND.R E107 870428))

((HEADING HEADING11416)
(SHIP.R HEADING11416 KNOX)
(SHIP.COURSE.R HEADING11416 195)
(POSITION.R HEADING11416
POSITION11410))

((READINESS READINESS11408)
(SHIP.R READINESS11408 KNOX)
(READINESS.LEVEL.R READINESS11408 C4))

((ARRIVE ARRIVE11400)
(SHIP.R ARRIVE11400 KNOX)
(TIME.R ARRIVE11400 870424)
(NEXT-ACTION.R ARRIVE11400 E107))

((POSITION POSITION11410)
(SHIP.R POSITION11410 KNOX)
(LONGITUDE.R POSITION11410 79)
(LATITUDE.R POSITION11410 18))
```

Hovy's RST Text Structurer

Initial goal:

```
(BMB SPEAKER HEARER
(SEQUENCE-OF E105 ?NEXT))
```

Simple gloss:

Tell the hearer the sequence of events of which E105 is a principal part.

More precisely:

Achieve the state in which the hearer believes that it is the intention of the speaker that they mutually believe that the event E105 is followed by some other event.

The resulting text:

Knox, which is C4, is en route to Sasebo. Knox, which is at 18N 79E, heads ssw. It arrives on 4/24. It loads for 4 days.

The Relation Definition for SEQUENCE

<i>relation name:</i>	SEQUENCE
<i>constraints on N:</i>	multi-nuclear
<i>constraints on the combination of nuclei:</i>	A succession relationship holds between the situations presented in the nuclei
<i>the effect:</i>	R recognises the succession relationships among the nuclei

The Relation/Plan for SEQUENCE

```
Name: SEQUENCE
Results: ((BMB SPEAKER HEARER (SEQUENCE-OF ?PART ?NEXT)))

Nucleus requirements/subgoals:
((BMB SPEAKER HEARER (TOPIC ?PART)))

Satellite requirements/subgoals:
((BMB SPEAKER HEARER (TOPIC ?NEXT)))

Nucleus+Satellite requirements/subgoals:
((NEXT-ACTION ?PART ?NEXT))

Nucleus growth points:
((BMB SPEAKER HEARER (CIRCUMSTANCE-OF ?PART ?CIR))
(BMB SPEAKER HEARER (ATTRIBUTE-OF ?PART ?VAL))
(BMB SPEAKER HEARER (PURPOSE-OF ?PART ?PURP)))

Satellite growth points:
((BMB SPEAKER HEARER (ATTRIBUTE-OF ?NEXT ?VAL))
(BMB SPEAKER HEARER (DETAILS-OF ?NEXT ?DETS))
(BMB SPEAKER HEARER (SEQUENCE-OF ?NEXT ?FOLL)))

Order: (NUCLEUS SATELLITE)
Relation-phrases: (" "then" "next")
Activation-question:
"Could ^A be presented as start-point, mid-point, or end-point of some succession of items along some dimension? -- that is, should the hearer know that ^A is part of a sequence?"
```

Hovy's RST Text Structurer

- nucleus and satellite requirements are treated as preconditions
- growth point goals are plan steps

Hovy's RST Text Structurer

The general algorithm:

Given

1. a goal
2. a set of input units

try to consume the set of input units:

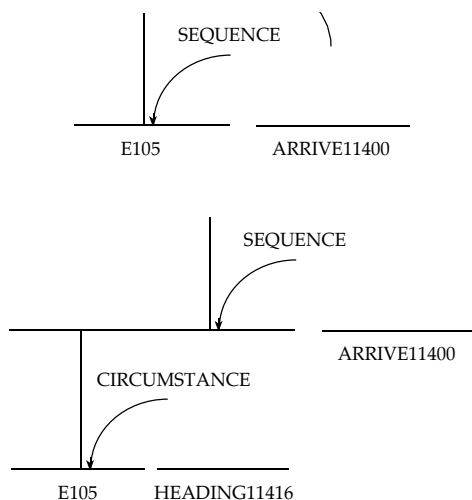
for each goal on the agenda

- find an operator that matches
- remove main topic from inputs
- add nucleus and satellite growth points to agenda

until the agenda is empty or there are no more input units

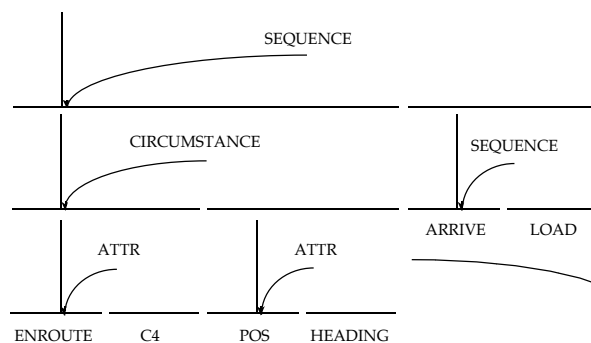
Hovy's RST Text Structurer

Steps in the construction of a paragraph:



Hovy's RST Text Structurer

The final text:



Knox, which is C4, is en route to Sasebo. It is at 79N 18E heading SSW. It will arrive on 4/24, and will load for four days.

Rhetorical Structure Theory

Advantages:

- RST relations tend to be smaller and less constraining than McKeown's schemata
- should result in a wider variety of texts and text structures, since smaller portions of the text can be varied by switching schemata

Rhetorical Structure Theory

Problems:

- this is only text organization—what about text selection?
- the domain dependence of semantics
- are we conflating intentional function and rhetorical function?
- just how many rhetorical schemata are there?

Do we want hybrid schemas and RST relations?

Some Issues

Questions for implementations of the theory:

- The relation definitions are underconstrained: we generate too many texts, and not all are good.
- Is rhetorical structure enough to guarantee coherence?

Questions about the theory itself:

- What is the taxonomy of relations? How many are there? What if we disagree on analyses?
- What is the grain-size of the smallest units?

Overview

1. Davey's PROTEUS
2. McKeown's TEXT
3. Kukich's ANA
4. RST in NLG
5. Wahlster *et al*'s WIP

The WIP Project

Aim:

to coordinate NL and graphics so they generate a multiplicative improvement in communication capabilities

Task:

multimodal presentation generation:

- multimodal explanations and instructions for assembling, using, maintaining or repairing physical devices
- German or English explanations for using espresso machine or assembling a lawn-mower

The WIP Project

Input:

a formal description of the communicative intent of a planned presentation


The modality choice question:

- what should go into text
- what should go into graphics
- what kind of links between verbal and non-verbal fragments are needed


WIP: An Example Output



Lift the lid.



To fill the watercontainer, remove the cover.



Use cold tapwater.

Generation Parameters

Output is customised for intended target audience and situation:

- user stereotypes (novice, expert)
- target languages (German, English)
- layout formats (hardcopy of instruction manual, screen display)
- output modes (incremental vs complete output only)

The Architecture of WIP

The Elements of WIP

Component technologies:

- terminological logics
- RST-based planning
- constraint processing techniques
- tree adjoining grammars with feature unification

Presentation Planning

Structure of a presentation is determined by a set of semantic and pragmatic relations:

- Semantic = domain structure: sequence, contrast, cause–result
- Pragmatic = communicative function of document parts: elaborate, illustrate

Knowledge for presentation planning consists of:

- presentation strategies
- revision strategies
- presentation metarules

Presentation Strategies

```
(def-presentation-strategy
  :Header (Describe P A
            (Orientation ?orientation) G)
  :Effect (BMB P A
            (Has-Orientation
             ?x ?orientation))
  :Applicability-Conditions
    (Bel P (Has-Orientation
            ?x ?orientation))
  :Main-Acts (S-Depict P A
              (Orientation ?orientation)
              ?p-ori ?picture)
  :Subsidiary-Acts
    (Achieve P
     (BMB P A
      (Identifiable A ?x ?px ?picture))
     ?mode))
```

Presentation Strategies

```
(def-presentation-strategy
:Header (Background P A ?x ?px ?picture G)
:Effect (BMB P A
         (Identifiable A
          ?x ?px ?picture))
:Applicability-Conditions
(AND
 (Bel P (Image-Of ?x ?px ?picture))
 (Bel P (Perceptually-Accessible A ?x))
 (Bel P (Part-Of ?x ?z)))
:Main-Acts (S-Depict P A
            (Object ?z) ?pz ?picture)
:Subsidiary-Acts
(Achieve P
 (BMB P A
  (Identifiable A ?z ?pz ?picture))
 ?mode))
```

Presentation Metarules

```
IF (IS-A ?current-attribute-value
    Spatial-Concept)
THEN (DoBefore *graphics-strategies*
      *text-strategies*)
```

Action Structure


Coherence

Three levels:


- syntactic (surface level phenomena, immediate connectivity between adjacent segments)
- semantic (content and global structuring of presentation: thematic organisation so perceived as a unified whole)
- pragmatic (effectiveness: has to be compatible with addressee's interpretative ability)

Syntactic Incoherence

Semantic Incoherence



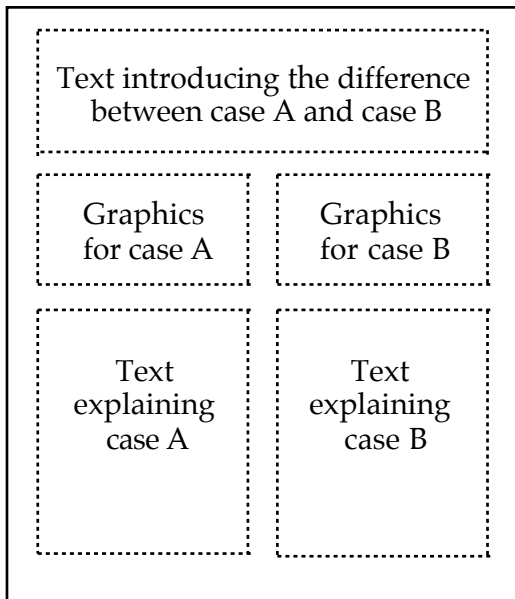
Multimodal Incoherence



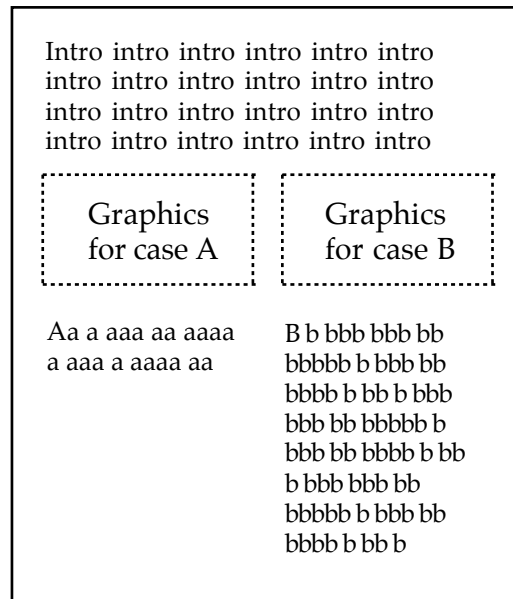
User Modelling

1. Espresso is coffee prepared in a special machine from finely ground coffee beans, through which steam under high pressure is forced. Because of the extreme pressure and high temperature, you should wait for at least two minutes after switching off the machine before you open the cover of the water container.
2. Wait at least for two minutes after switching off the machine before you open the cover of the water container.

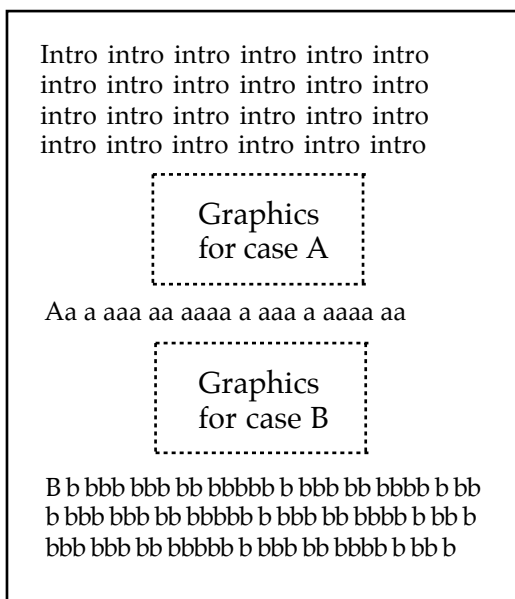
Coordinating Content and Layout



Coordinating Content and Layout



Coordinating Content and Layout



WIP: Summary

Basic ideas:

- common representation as input leads to media choice and media interaction
- ideas and concepts from NLG generalise to communication in general
- design is a non-monotonic planning process

Prototype fully implemented: 30 person years effort.

Related Work

- XTRA [Allgayer *et al*1989]
- CUBRICON [Neal and Shapiro 1988]
- ALFRESCO [Stock 1991]
- SAGE [Roth *et al*1988]
- FN/ANDD [Marks and Reiter 1990]
- COMET [Feiner and McKeown 1989]

Other Issues in Discourse Organization

Recent work that suggests interesting new directions:

- Rambow's work on domain communication knowledge
- Sibun's work on the local organization of text
- Mooney, Carberry and McCoy's work on the basic block model

Summary

What we've seen:

- McKeown's use of text schemata: limited but computationally attractive
- Rhetorical Structure Theory: theoretically more appealing than schemas, more flexible, but throws up more unanswered questions
- WIP: multimodal document generation

What's Coming Next ...

1. An Overview of NLG
2. Linguistic Realization
3. Text Planning
4. Generating Referring Expressions