Microplanning

Albert Gatt
Institute of Linguistics, University of Malta
http://staff.um.edu.mt/albert.gatt/
albert.gatt@um.edu.mt

The “consensus” architecture

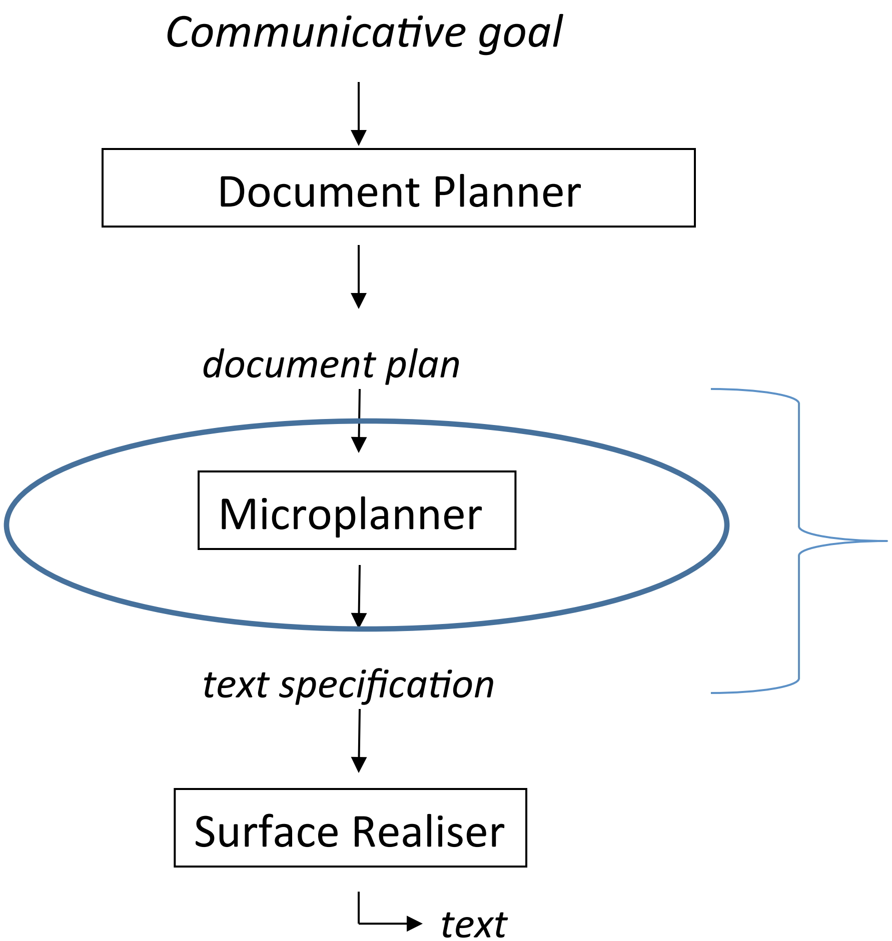

Focus today:

• Document plan includes the messages/ events to mention.
• Microplanner has to flesh out the “linguistic” content of those messages/events.

The “consensus” architecture

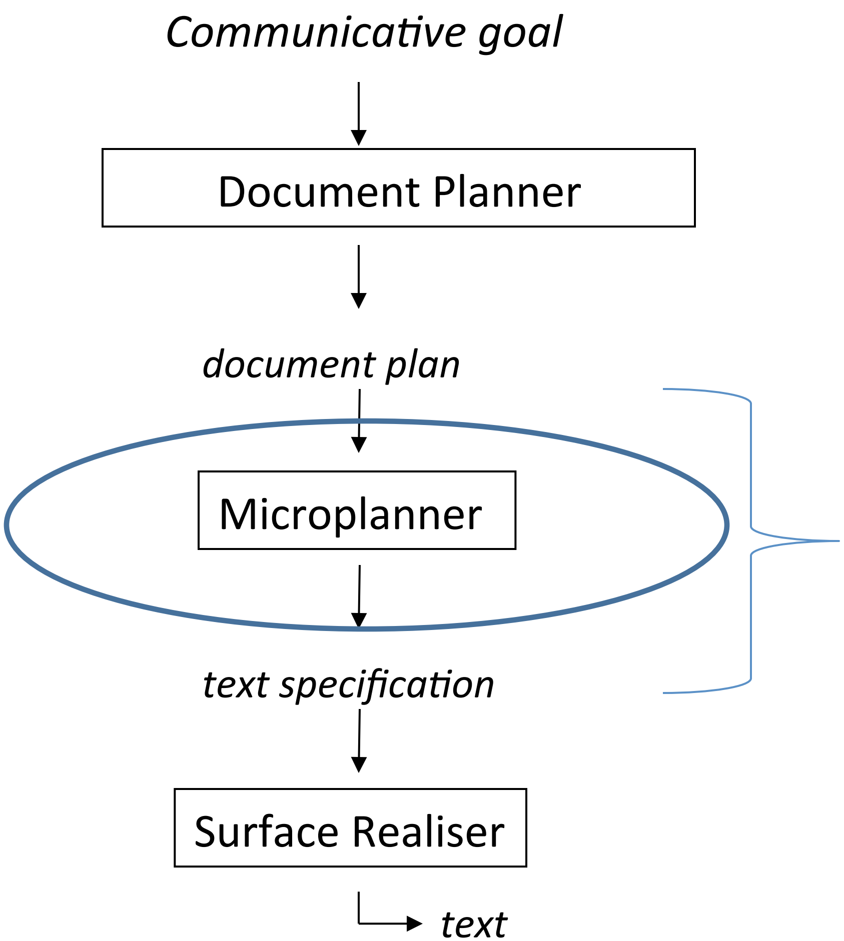

What happens in microplanning?
- Lexicalisation: choosing words (but not only)
- Referring expressions (descriptions, pronouns)
- Aggregation

Preview

1. Intro case study: BT-Nurse micro-example
2. Lexicalisation
3. Aggregation
4. Referring expression generation
   - Determining form: Using context and salience
   - Classic algorithms for definite descriptions

Case study: BabyTalk (BT-Nurse)

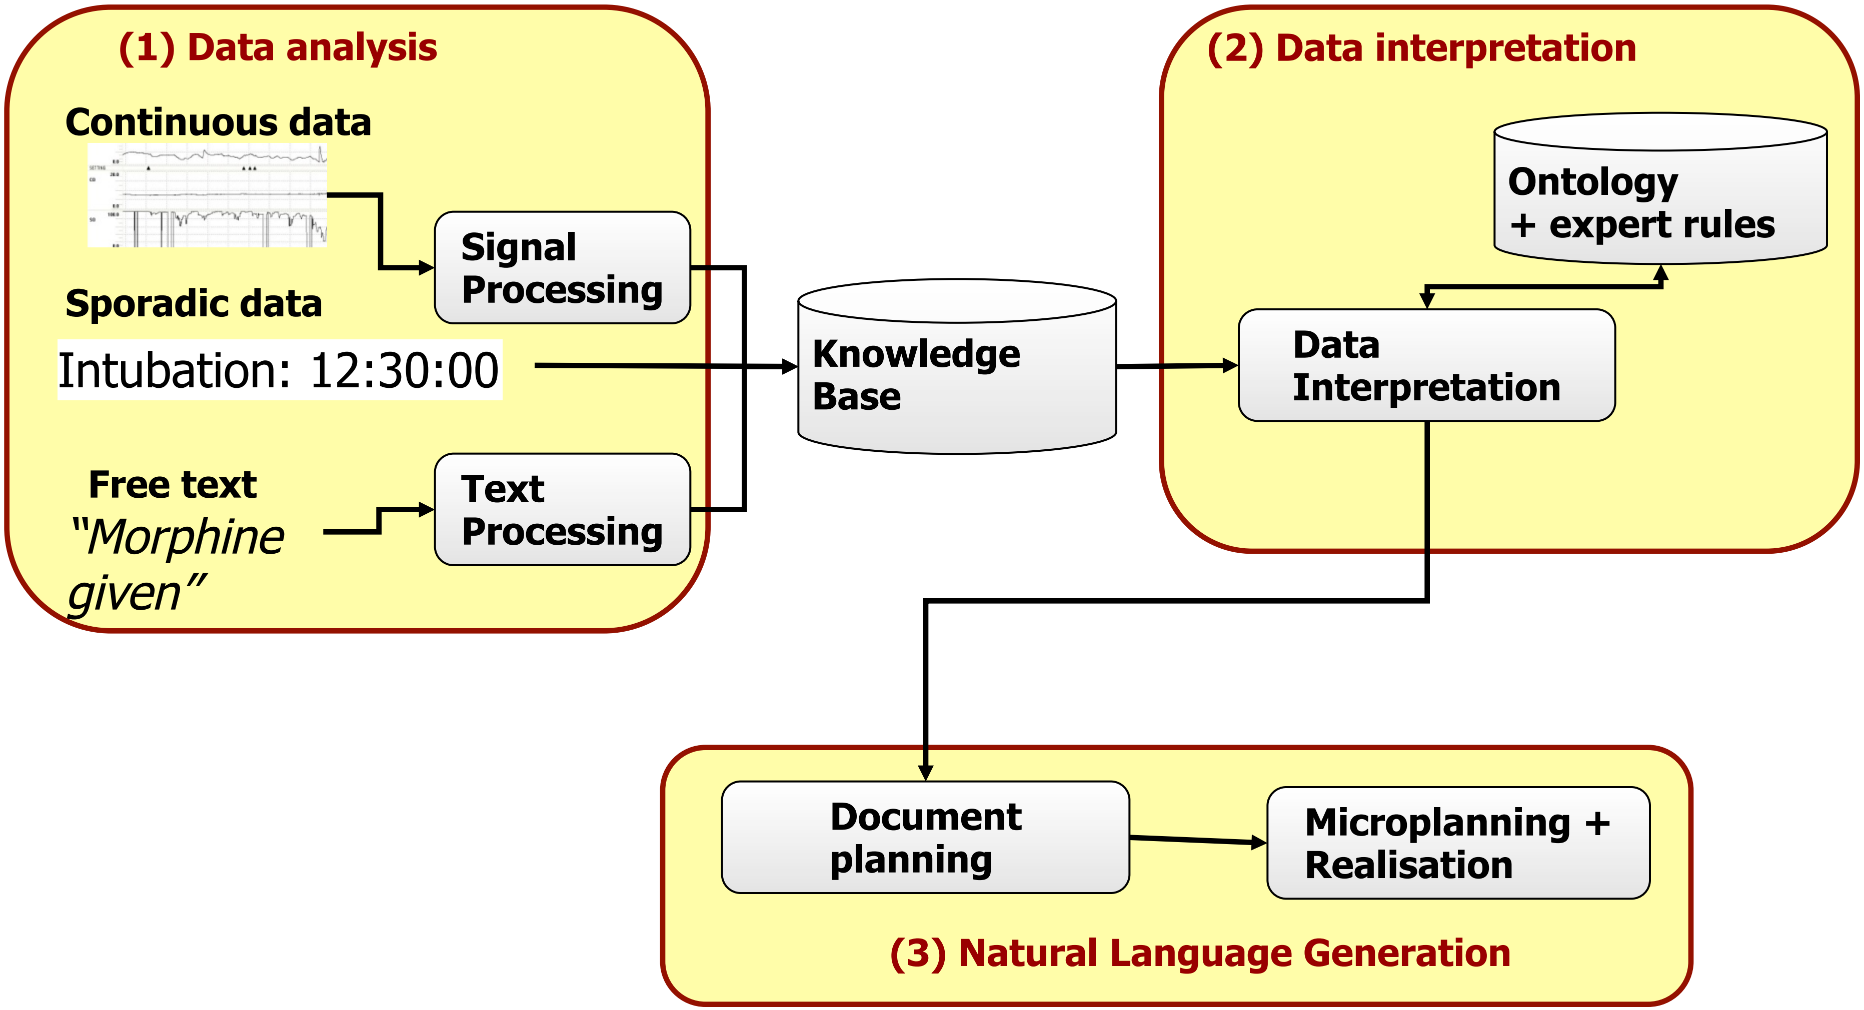

Gatt et al, 2009

A micro example

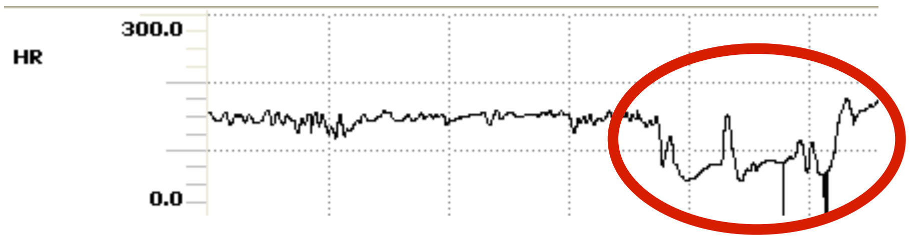

Input data: unstructured raw numeric signal from patient's heart rate monitor (ECG)

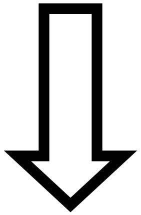

There were 3 successive bradycardias down to 69.

A micro example: pre-NLG steps

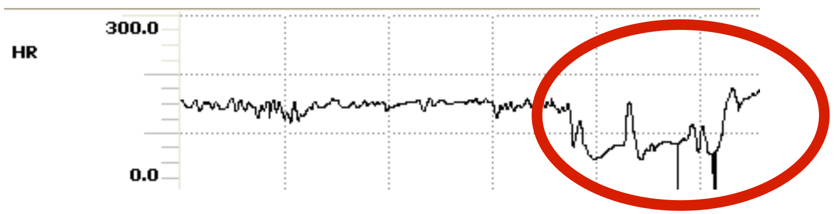

SEQUENCE (BRADYCARDIA)

BRADYCARDIA (16:58:46) Imp: 31.64
BRADYCARDIA (17:01:15) Imp: 79.80
BRADYCARDIA (17:03:57) Imp: 80.21
BRADYCARDIA (17:04:30) Imp: 39.97
BRADYCARDIA (17:05:01) Imp: 34.60
BRADYCARDIA (17:06:03) Imp: 66.24

(1) Signal Analysis (pre-NLG)

- Identify interesting patterns in the data.
- Remove noise.

(2) Data interpretation (pre-NLG)

- Estimate the importance of events
- Perform linking & abstraction

A micro example: Document planning

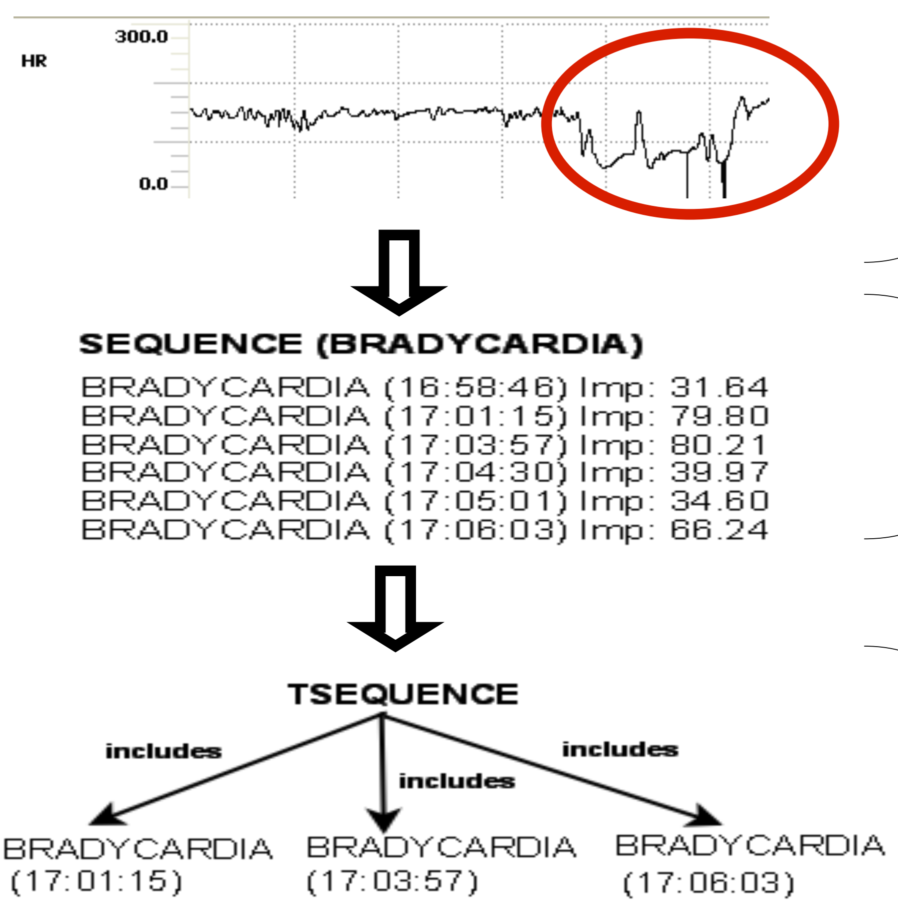

(1) Signal Analysis (pre-NLG)

- Identify interesting patterns in the data.
- Remove noise.

(2) Data interpretation (pre-NLG)

- Estimate the importance of events
- Perform linking & abstraction

(3) Document planning

- Select content based on importance
- Structure document using rhetorical relations
- Communicative goals (here: *assert something*)

A micro example

$$\begin{bmatrix} \textit{Event} & & \\ \textbf{TYPE} & \textit{existential} & \\ \textbf{PRED} & \textit{be} & \\ \textbf{TENSE} & \textit{past} & \\ \textbf{ARGS} & \begin{bmatrix} \textbf{THEME} & br_{01} \\ \textbf{VALUE} & 69 \end{bmatrix} & \end{bmatrix}$$

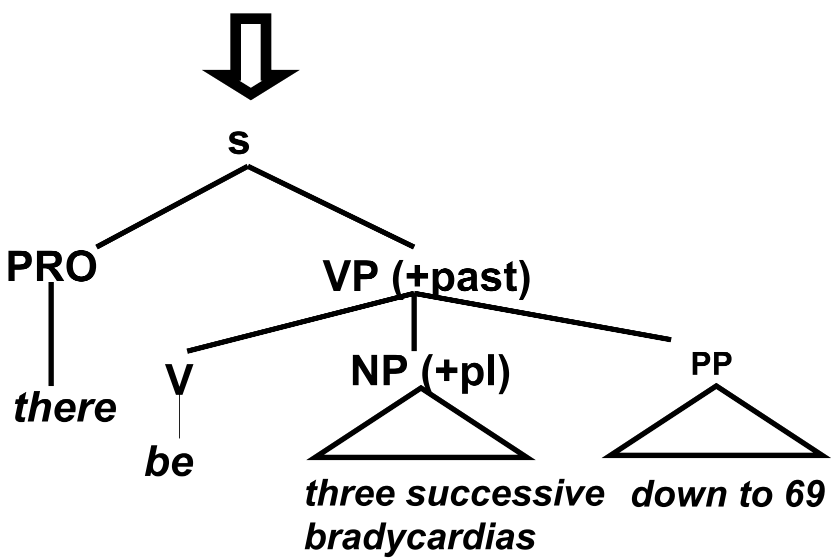

(4) Microplanning

Map events to semantic representation

- lexicalise: *bradycardia* vs *sudden drop in HR*
- aggregate multiple messages (3 bradycardias = one sequence)
- decide on how to refer (*bradycardia* vs *it*)
- choose sentence form (*there were…*)
- Referring expressions

(5) Realisation

- map semantic representations to syntactic structures
- apply word formation rules

Part 1

LEXICALISATION

Choosing words

- In many cases, done in a rule-based fashion.

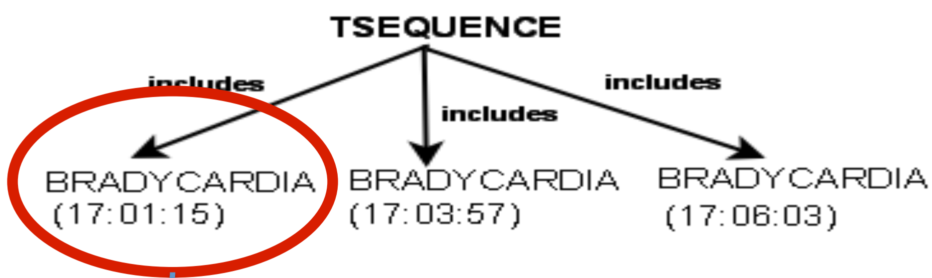

Lexical rule base

words

Choosing words

- Actually, it's very rarely just about individual words.
  - "world" → language
    - Not a straightforward mapping

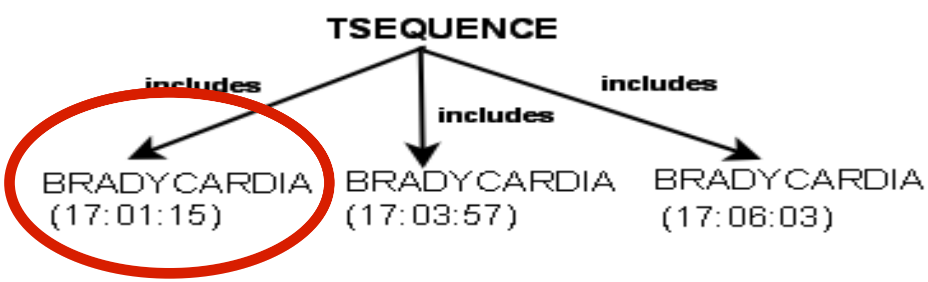

- Is this a thing (to be expressed as a NP)?
- Is this an event?

Choosing words

- In many cases, done in a rule-based fashion.

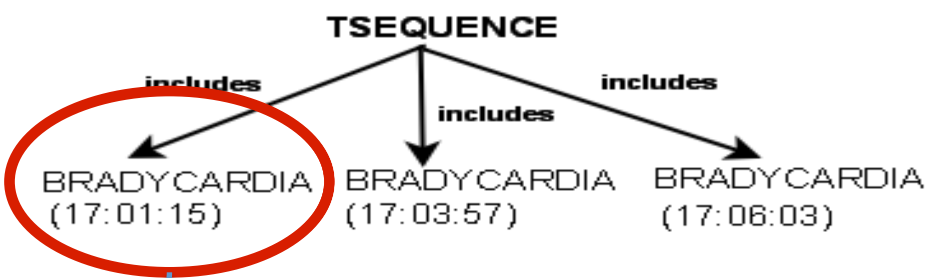

Lexical rule base

- Verbnet
- Framenet
- Wordnet

If X is-a *bradycardia*, then:

Choose verb: *be*
Make X the **theme**
Make X's max the **value**

Construction: ***existential***

$$\left[\begin{array}{ll} \textit{Event} & \\ \textbf{TYPE} & \textit{existential} \\ \textbf{PRED} & \textit{be} \\ \textbf{TENSE} & \textit{past} \\ \textbf{ARGS} & \left[\begin{array}{ll} \textbf{THEME} & \textit{br}_{01} \\ \textbf{VALUE} & 69 \end{array}\right] \end{array}\right]$$

Just words or also syntax?

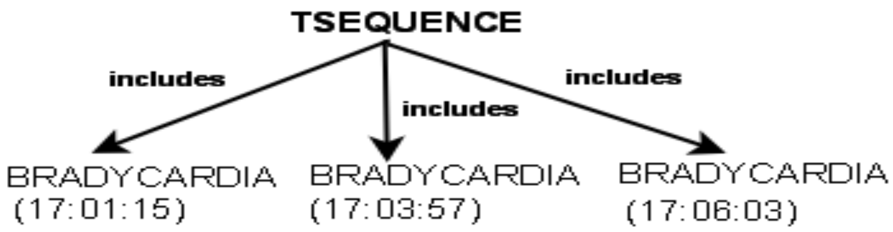

- Here, we just have an ontology instance, which specifies that there was a particular event, at a certain time.
- Choice of verb makes a difference to argument structure and to syntax.
  - *There **was** a bradycardia down to 69.*
  - *The bradycardia **went down** to 69.*

Just words or also "focus"?

- Word choice has consequences for how parts of a concept are "packaged".
  - The event, the manner, the time…
  - *HR* ***rose*** *to YYY* ***over the next 5 minutes.***
    - Verb describes direction of "motion".
    - Time is lexicalised as a PP.
  - *HR* ***shot up*** *to YYY over the next five minutes*.
    - Verb describes manner and direction.
    - Time still lexicalised as a PP.
  - *HR* ***reached*** *YYY over the next five minutes.*
    - Verb incorporates time.
    - Direction is left implicit.

Genre conventions

- Lexicalisation also depends on conventional ways of describing things.

KB Instance:
CMV
O2 = 23

KB Instance:
CMV
O2 = 21

- Medics always say: ***X is on CMV***
  - *...in 23% oxygen*
  - *...in air*

User/audience modelling

- Who are we generating for?
  - E.g. expert/non-expert
  - E.g. parent/guardian

- Example (Mahamood and Reiter 2011):
  - *the baby was put on HFOV*
  - *your child was put on a High Frequency Oxygen Ventilator (HFOV) to aid her breathing.*

User/audience modelling

- Janarthanam and Lemon (2014):
  - Dialogue system: instructions to user
  - *Please plug in the broadband cable*
  - *Please plug in the thin white cable with grey ends*
- Approach based on Reinforcement Learning:
  - Corpus of dialogues, annotated with success measures.
  - Learning: policy to maximise success, depending on user expertise.

Stylistic variation

- Lexical and syntactic choice to convey:
  - Degrees of formality (e.g. Paiva & Evans ’05)
  - Personality (e.g. Mairesse & Walker ‘11)
  - Affect (e.g. Mahamood & Reiter ‘11)

- Mahamood & Reiter:
  - Aim: minimise stressfulness for guardians of sick patients.
  - Method:
    - compute a predicted stress level for a message
    - High stress → select mitigating expressions

Since last week, his inspired Oxygen (FiO2) was lowered from 56% to 21% (which is the same as normal air). This is a positive development for your child.

So what are the choices?

- Paradigmatic choices:
  - Given an input concept or conceptual structure, choose from among a potentially wide range of possible words for (parts of) the input.
- Syntagmatic choices:
  - Choices will have structural consequences for how the input is mapped to a syntactic structure, and eventually to a linearised string.

Syntagmatic & paradigmatic choices

Syntagmatic choice: verb frame

There **be** **THEME** to **VALUE**
There was a bradycardia

THEME **occur** reaching **VALUE**]
There was a bradycardia

Paradigmatic choice: verb

be ***occur***

Considerations:
- Conventions
- Style, affect
- Focus/foregrounding

Input:

BRADYCARDIA
(17:01:15)

Words or syntax? Elhadad & McKeown (1997)

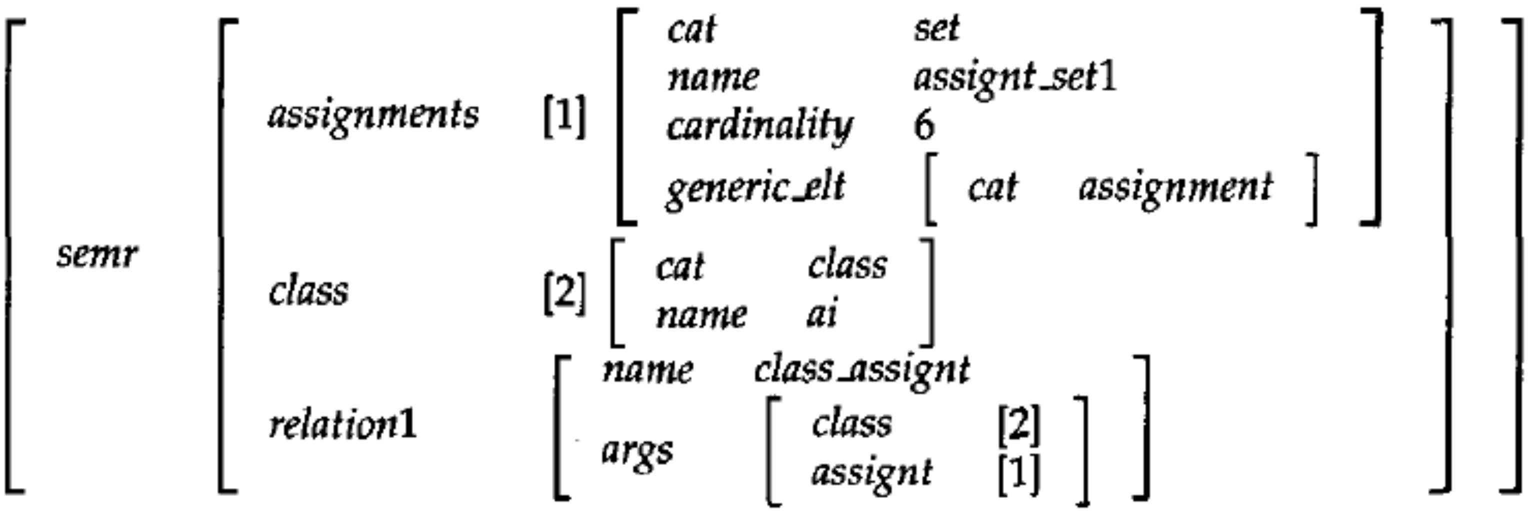

- Input semantics
  - Domain: University education
  - A feature structure specifying the number of assignments due in a given class.
- Lexicaliser performs a search in a library of structures, finding the one that can be unified with the input.

Words or syntax? Elhadad and McKeown (1997)

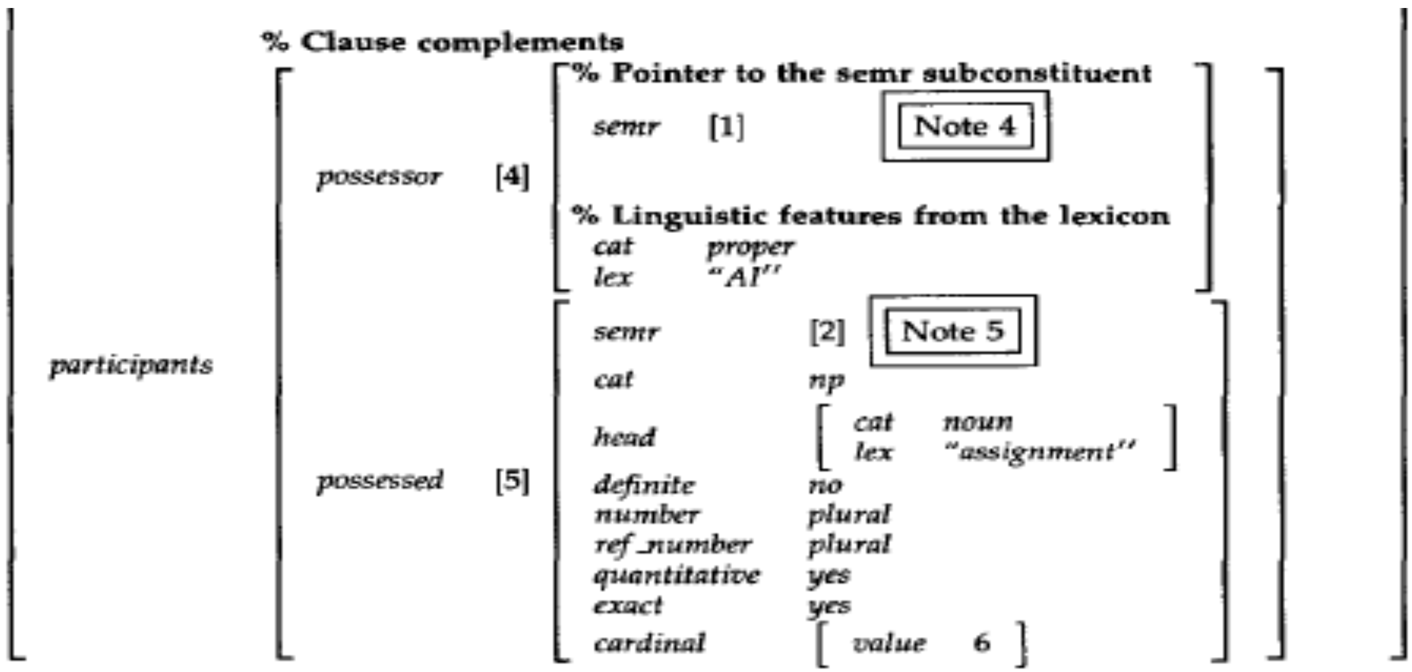

- Part of the output:
  - Syntactic specification
  - Word choice

Words or syntax? Statistical approaches.

Events:	skyCover_1	temperature_1				windDir_1		windSpeed_1	
Fields:	*percent*=0-25		*time*=6am-9pm	*min*=9	*max*=21	*mode*=S			mean=20
Text:	cloudy ,	with	temperatures between	10	20 degrees .	south	wind	around	20 mph .

- Lexicalisation = "alignment" between data and word sequences.
  - Liang et al (2009); Konstas & Lapata (2013)
  - DB consisting of records (r), with field-values (f, v) and a type (t)
  - Alignment: $p(\mathbf{w} \mid r, f, t) = \prod_{j}^{|w|} p(w_j \mid r, t, f, v)$

Lexicalisation: Interim summary

- Trivial case:
  - One piece of input → one word
- Non-trivial (most frequent) case:
  - One piece of input → several lexical choices
  - Choice has consequences for:
    - Syntactic structure
    - Foregrounding/emphasis
    - Style, affect

Part 2

AGGREGATION

Why aggregate?

- Aggregation makes text more consise, fluid, readable (Dalianis 1999; Cheng 2000)

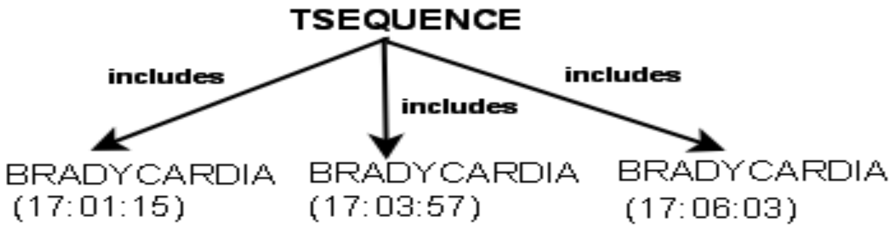

- Three events, all of the same kind.
  - There was a bradycardia at 17:01 down to ...
  - There was a bradycardia at 17:03 down to...
  - There was a bradycardia at 17:06 down to...

Where does it happen?

“Semantic” aggregation

- Merge based on semantic information.
- Can be done at microplanning level.

Syntactic aggregation

- Realise messages.
- Merge based on phrase structure.
- See, e.g. Harbusch & Kempen 2009

Semantic aggregation in BabyTalk

- Aggregation rules fire every time a new event is added to the DM.
- Rules are sensitive to event type, discourse relation, …
  - E.g. are two events in a causal relation?
  - Are two events of the same type and close together in time?
- If preconditions satisfied, merge semantic structures.

Example 1

Aggregation rule:

If e IS-A BRADYCARDIA then:

- check if another bradycardia is in DM
- if they are in a TSEQUENCE then:
    - Merge THEME args
    - Choose lowest VALUE
    - (Merge start/end times)

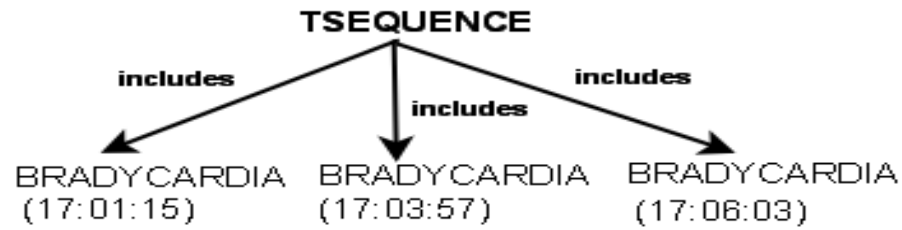

br_{01} br_{02} br_{03}

$$\begin{bmatrix} \textit{Event} & \\ \textbf{TYPE} & \textit{existential} \\ \textbf{PRED} & \textit{be} \\ \textbf{TENSE} & \textit{past} \\ \textbf{ARGS} & \begin{bmatrix} \textbf{THEME} & br_{01} \\ \textbf{VALUE} & 70 \end{bmatrix} \end{bmatrix} + \begin{bmatrix} \textit{Event} & \\ \textbf{TYPE} & \textit{existential} \\ \textbf{PRED} & \textit{be} \\ \textbf{TENSE} & \textit{past} \\ \textbf{ARGS} & \begin{bmatrix} \textbf{THEME} & br_{01} \\ \textbf{VALUE} & 69 \end{bmatrix} \end{bmatrix} \Rightarrow \begin{bmatrix} \textit{Event} & \\ \textbf{TYPE} & \textit{existential} \\ \textbf{PRED} & \textit{be} \\ \textbf{TENSE} & \textit{past} \\ \textbf{ARGS} & \begin{bmatrix} \textbf{THEME} & \{br_{01}, br_{02}\} \\ \textbf{VALUE} & 69 \end{bmatrix} \end{bmatrix}$$

Some other examples

- Merging related trends into one “complex” event:
  - *At around 23:30, urine output rate rose […] and had dropped by around 05:15.*
- Merging causally connected events:
  - *The baby was intubated and was put on CMV.*
  - Note: causality is left implicit.

Challenges in aggregation

- Control:
  - How many events to merge? At what point does aggregation become "too much"?
    - Use "sensible" thresholds?
    - Probably quite domain-dependent.
  - Under what conditions should aggregation be performed?
    - E.g. merge events which are very far apart?

- Explicit/implicit connectives:
  - Which relations should be made explicit?
    - *X caused Y*… vs *X and Y*
    - Some relations can easily be inferred.

Is aggregation always useful?

- The typical answer is:
  - Yes, aggregation makes text more readable, less repetitive. (e.g. Reape & Mellish 1999)
- De Rosis & Grasso (2000; cf. Walker 1997):
    - How difficult is the topic for the reader?
    - How stressful will it be?
    - How deserving is it of emphasis?
  - Aggregation can be used for strategic purposes:
    - Aggregate negative messages.
    - Emphasise positive messages.

Using aggregation strategically

Output 1

However, I must inform you that this drug may cause some side effects. The first one is nausea; it is serious, it occurs infrequently, in a strong form, in sensitive patients. The second one is headache; it is serious, it occurs infrequently, in a strong form, in sensitive patients. The third one is insomnia; it is not serious, it occurs frequently, in a strong form, in sensitive patients.

Output 2 (after plan revision)

However, I must inform you that this drug may cause some side effects. **A first group of them includes nausea [...] and headache [...]; these side effects are both serious.** Then, you may have insomnia: it is not serious but can be frequent; however, once again I would like to reassure you that it occurs only in particularly sensitive patients. All these side effects can occur in a strong form.

De Rosis and Grasso (2000)

Doing it statistically: Walker et al 2002

You are leaving from Newark.
You are leaving at 5.

MERGE

You are leaving from Newark at 5.

You are leaving from Newark.
You are going to Dallas.

CONJUNCTION

You are leaving from Newark and going to Dallas.

- Finite set of merge/aggregation rules, operating on sentence plans.
- Train a boosting algorithm:
  - Sample of sentence plans, annotated with pragmatic and syntactic information.
  - Ratings by human judges: *in the context of the dialogue, is this a good sentence?*
- Generation:
  - Given a text plan, generate several sentence plans.
  - Boosting function ranks them (what is the predicted human preference?)
  - Output the best result.

Doing it statistically: Barzilay & Lapata 2006

Passing					
PLAYER	CP/AT	**YDS**	**AVG**	**TD**	INT
Cundiff	22/37	237	6.4	1	1
Carter	23/47	237	5.0	1	4
...	...	...	...	...	...

- Aggregation as supervised partitioning:
  - Aggregate DB entries depending on how **similar** they are.
  - Input: DB entries chosen by the content planner.
  - Output: a partition of the set, s.t. every element of the input occurs in exactly one subset.
    - If 2 entries are in a subset, then they are to be aggregated.

Doing it statistically: Barzilay & Lapata 2006

Passing					
PLAYER	CP/AT	**YDS**	**AVG**	**TD**	INT
Cundiff	22/37	237	6.4	1	1
Carter	23/47	237	5.0	1	4
...	...	...	...	...	...

- Binary classification?
  - For any 2 entries, classify as *aggregate* or *not*.
  - Could do, but can't really deal with global constraints:
    - Don't aggregate more than 3 times in total…
    - If you aggregate A and B, you need to aggregate C as well…
- Global optimisation:
  - Use pairwise classifier to compute probability of aggregation for pairs, based on similarity.
  - Find a global partitioning which:
    - Maximises the sum of pairwise scores
    - Respects the constraints

Aggregation: Interim Summary

- Probably, highly domain-dependent.
- Rule-based approaches (e.g. BabyTalk) perform quite well, but very labour-intensive.
  - Advantage: full control.
- Data-driven approaches seem to do well, but on quite restricted (structured) domains (DB input).

- Open question:
  - When is aggregation appropriate?
  - How much aggregation is appropriate?

Part 3

REFERRING EXPRESSIONS

What is a referring expression?

Any expression which serves to identify any thing, process, event, action, or any other kind of individual or particular I shall call a referring expression. Referring expressions point to particular things; they answer the questions Who?, What?, Which?

(John Searle (1969). *Speech Acts: An Essay in the Philosophy of Language*. Cambridge: CUP)

What is a referring expression?

- Typically, a noun phrase.

- Two important choices:
  - What form should the RE take?
  - What content should be chosen for the RE?

Zooming in

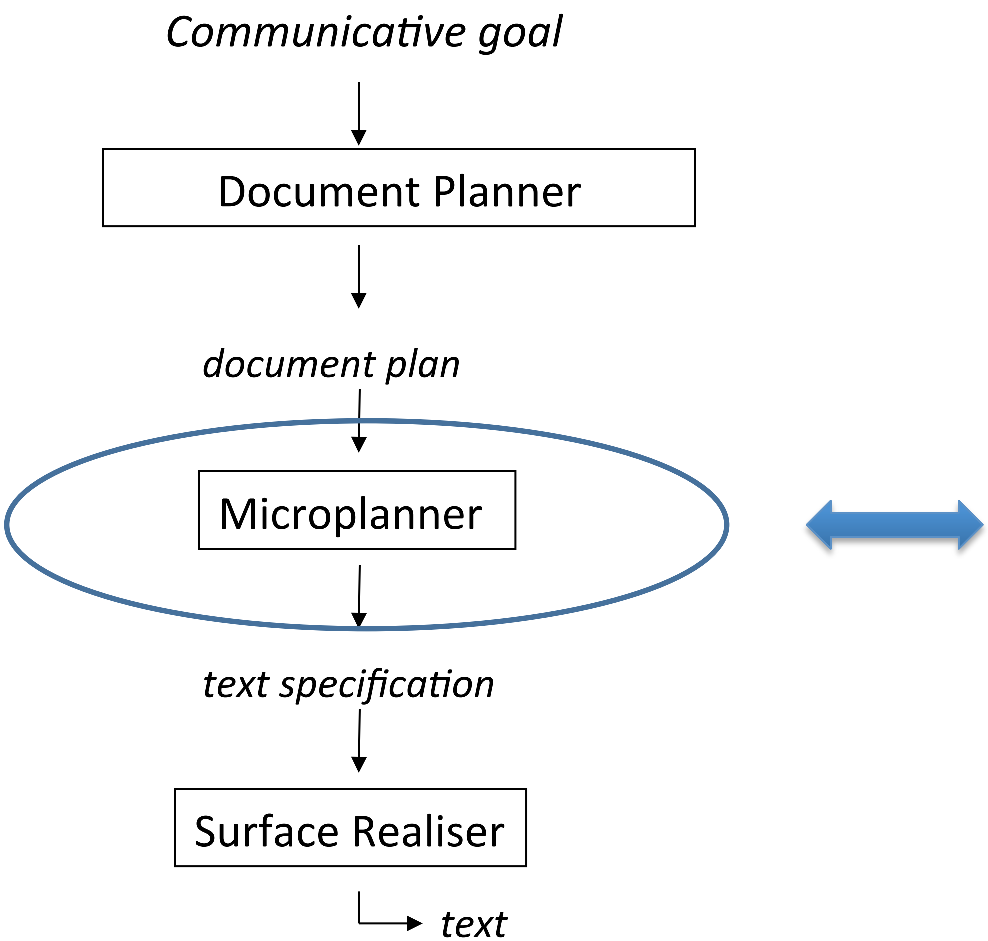

Discourse model

- Record of entities mentioned so far

+

- Syntactic (or thematic?) role of entities.

Determining form

- We need some framework to account for the **salience** of discourse entities, e.g.:
  - Centering Theory (Grosz et al 1995)
    - Salience primarily determined by syntactic role (Subject > Object > Other)
    - Aim is to maintain transitions between messages as smoothly as possible.
  - Accessibility Theory (Ariel 2001)
    - Different types of NPs signal to the hearer the degree to which the entity in question is accessible.
- NB: Some of these frameworks assume that **syntactic** info is available!

Two strategies

Discourse Model

The patient was put on HFOV.

The percentage O2 is 21%.

Next message:

The patient/She was intubated.

Strategy 1 (easy):

- Has the entity been mentioned in the previous sentence?
  - If yes, use pronoun
  - If not, use definite description.
- Cf. Dale, 1992

Two strategies

Discourse Model

The patient was put on HFOV.

The percentage O2 is 21%.

Next message:

The patient/She was intubated.

Strategy 2 (more elaborate)

- Compute the salience of the entity
  - Grammatical role? Information structure? Recency?
  - Salience is dynamic. Function must decay!
- If the entity is the most salient of its type, use a pronoun.
- Else, use a (possibly reduced) definite description.

- Cf. Krahmer & Theune 2001; McCoy & Strube 1999; Callaway & Lester 2002, …

Open questions

- Computing salience requires syntactic information in some frameworks.
  - Is the microplanning/realisation distinction tenable?

- Use of pronouns has an impact on discourse coherence, e.g. in Centering Theory.
  - Sometimes, a pronoun can cause a sudden “shift” in the thread of discourse.
  - Is microplanning really separate from document planning?
  - Kibble & Power 2004: plan text to maximise coherence and ease of pronoun resolution for the reader.

Selecting content for definite descriptions

$$\begin{bmatrix} \textit{Event} & \\ \textbf{TYPE} & \textit{existential} \\ \textbf{PRED} & \textit{be} \\ \textbf{TENSE} & \textit{past} \\ \textbf{ARGS} & \begin{bmatrix} \textbf{THEME} & br_{01} \\ \textbf{VALUE} & 69 \end{bmatrix} \end{bmatrix}$$

We need to decide what to say about this “entity”.
Not the only bradycardia in the KB.
Not everything we know about it may be relevant.

Things we might know about *br01:*

- type (= bradycardia)
- time of occurrence
- minimum value
- maximum value
- …

Another example: visual domain

Our KB

Our message

- Suppose our document planner has included this message:
  - **bottom-right(E3)**
- We want to say something like:
  - "E3 is in the bottom right"
- But our user doesn't know what E3 is.
- We need to **describe** it.

REG is a search problem!

The KB as we see it

Given:

A KB with objects and properties

A target referent

Find:

A combination of properties that will distinguish the target referent from its distractors

Problem definition

How would you distinguish the object in the red box?

Overspecified:

the red chair facing back

the large red chair

Underspecified:

the chair

the red chair

Minimally specified:

the large chair

the chair facing back

KB + referent

Distinguishing description

Defining adequacy 1: Gricean

Gricean Maxim of Quantity (Grice 1975): Say no more than you must

- Produce the briefest possible description
- (I.e. search through all possible descriptions in order of increasing length, until you find one that is distinguishing.)
- Problems:
  - Inefficient
  - Not necessarily "humanlike"

An incremental framework

Input: KB + target referent

1. Start by initialising an empty description D

2. **while** D does not distinguish the referent **do:**
   1. P $\leftarrow$ next property of the target referent to consider
   2. **if** the property excludes some distractors, **then**:

      remove the distractors

      add property to D

3. **return** the description

The important question

- Our "general" framework said: *pull out the next property* of the target.
  - How do we determine which one?
  - This is where, in our "general" algorithm, we need to factor in our definition of adequacy.
- Our strategy should:
  - Maximise the likelihood that we end up with a "good" description.
  - Depending on how we define "good", we might want to define how the next property to consider is identified.

Back to Grice

- Suppose we stick to the idea that "short is good".
  - Generating a minimal description is intractable.
  - BUT we can try to approximate this in our incremental framework.

- Greedy algorithm (Dale 1989):
  - Pull out the next property that removes the highest number of distractors.

The Greedy Algorithm

	type	colour	Size	orientation
E1	Chair	Black	Large	Front
E2	Chair	Red	Small	Front
E3	Chair	Red	Large	Back
E4	Sofa	Green	large	right

The KB

- Input: KB + target referent

1. Start by initialising an empty description D
2. **while** D does not distinguish the referent **do**
    1. Get the **most discriminatory** property of the target.
    2. **if** the property excludes some distractors, **then:**

       remove the distractors

       add property to D
3. **return** the description

The psycholinguistic evidence

- People seem to overspecify (contra Grice?).
- Speech production is **incremental**.
  - We don't compare all possible descriptions and choose the "best " (most Gricean) one.
  - We construct the description piece by piece, adding properties as we go along.
  - Not all properties are equal
    - Some properties seem to be used even when not required.
  - Can we do something similar computationally?

The Incremental Algorithm

- Algorithm proposed by Dale and Reiter (1995).
- Models REG as a search problem where:
  - A description is built piece by piece.
  - Some properties are given priority (they are tried first).
- The core element is a **preference order** which determines which properties will be tried out first.
  - **E.g. TYPE > COLOUR > ORIENTATION > SIZE**

The Incremental Algorithm

	type	colour	Size	orientation
E1	Chair	Black	Large	Front
E2	Chair	Red	Small	Front
E3	Chair	Red	Large	Back
E4	Sofa	Green	large	right

The KB

- Input: KB + target referent
- Input 2: preference order

1. Start by initialising an empty description D

2. **while** D does not distinguish the referent **do**
   1. Find the next attribute on the preference order
   2. Get the property for the target
   3. **if** the property excludes some distractors, **then:**
      remove the distractors
      add property to D

3. **return** the description

The Incremental Algorithm: example

	type	colour	Size	orientation
E1	Chair	Black	Large	Front
E2	Chair	Red	Small	Front
E3	Chair	Red	Large	Back
E4	Sofa	Green	large	right

The KB

- Preference order:

  type > colour > orientation > size

- D ← {}

The Incremental Algorithm: example

	type	colour	Size	orientation
E1	Chair	Black	Large	Front
E2	Chair	Red	Small	Front
E3	Chair	Red	Large	Back
~~E4~~	~~Sofa~~	~~Green~~	~~large~~	~~right~~

The KB

- Preference order:

  type > colour > orientation > size
- Next property: *Chair*
- Excludes *E4*
- *D ← {Chair}*

The Incremental Algorithm: example

	type	colour	Size	orientation
~~E1~~	~~Chair~~	~~Black~~	~~Large~~	~~Front~~
E2	Chair	Red	Small	Front
E3	Chair	Red	Large	Back
~~E4~~	~~Sofa~~	~~Green~~	~~large~~	~~right~~

The KB

- Preference order:

  type > colour > orientation > size

- Next property: *Red*
- Excludes *E1*
- *D* ← *{Chair, Red}*

The Incremental Algorithm: example

	type	colour	Size	orientation
~~E1~~	~~Chair~~	~~Black~~	~~Large~~	~~Front~~
~~E2~~	~~Chair~~	~~Red~~	~~Small~~	~~Front~~
E3	Chair	Red	Large	Back
~~E4~~	~~Sofa~~	~~Green~~	~~large~~	~~right~~

The KB

- Preference order:

  type > colour > orientation > size

- Next property: *Back*
- Excludes *E2*
- *D* ← *{Chair, Red, Back}*

The Incremental Algorithm: example

	type	colour	Size	orientation
~~E1~~	~~Chair~~	~~Black~~	~~Large~~	~~Front~~
~~E2~~	~~Chair~~	~~Red~~	~~Small~~	~~Front~~
E3	Chair	Red	Large	Back
~~E4~~	~~Sofa~~	~~Green~~	~~large~~	~~right~~

The KB

- The outcome is the description {chair, red, back}
- Could be realised as "the red chair facing backwards"
- Observe that this description is overspecified:
  - We didn't really need COLOUR at all! ORIENTATION alone would have done the trick.
  - Overspecification is an automatic consequence of the algorithm.

Summary on referring expressions

- Form:
  - Many competing frameworks to determine salience.
  - No single, agreed perspective in NLG.
  - Depending on framework, may turn out not to be a "pure" microplanning problem.

- Content:
  - REG has become a topic of research in its own right.
  - Several developments on the "classic" algorithms (see van Deemter & Krahmer, 2012)
  - An area in which psycholinguistic and computational work often inform each other!
    - More on this with Kees van Deemter on Friday morning.

Summary on microplanning

- Important open questions:
  - How separable is it from other NLG tasks?
    - Realisation (aggregation, lexicalisation)
    - Text planning (coherence, referential form)

- Current trend:
  - Data-driven approaches
  - Reliance on data-text alignment
  - Blurring of the boundaries between microplanning and other modules.

References

- Ariel, M. (2014). Accessing Noun Phrase Antecedents. London: Routledge.
- Barzilay, R., & Lapata, M. (2006). Aggregation via Set Partitioning for Natural Language Generation. In Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL'06) (pp. 359–366). New York: Association for Computational Linguistics.
- Callaway, C. B., & Lester, J. C. (2002). Narrative prose generation. Artificial Intelligence, 139(2), 213–252. doi:10.1016/S0004-3702(02)00230-8
- Cheng, H., & Mellish, C. (2000). Capturing the interaction between aggregation and text planning in two generation systems. In Proceedings of the first international conference on Natural language generation (INLG '00) (Vol. 14, pp. 186–193). Morristown, NJ, USA: Association for Computational Linguistics. doi:10.3115/1118253.1118279
- Dale, R. (1989). Cooking up referring expressions. In Proceedings of the 27th annual meeting on Association for Computational Linguistics (ACL'89) (pp. 68–75). Vancouver, BC: Association for Computational Linguistics. doi:10.3115/981623.981632
- Dale, R., & Reiter, E. (1995). Computational Interpretations of the Gricean Maxims in the Generation of Referring Expressions. Cognitive Science, 19(2), 233–263. doi:10.1207/s15516709cog1902_3
- Dalianis, H. (1999). Aggregation in Natural Language Generation. Computational Intelligence, 15(4), 384–414.
- De Rosis, F., & Grasso, F. (2000). Affective Natural Language Generation. In A. M. Paiva (Ed.), Affective interactions (pp. 204–218). Berlin and Heidelberg: Springer. doi:10.1007/10720296_15
- Elhadad, M., Robin, J., & McKeown, K. (1997). Floating constraints in lexical choice. Computational Linguistics, 23(2), 195–239. Retrieved from http://dl.acm.org/citation.cfm?id=972696

References

- Gatt, A., Portet, F., Reiter, E., Hunter, J., Mahamood, S., Moncur, W., & Sripada, S. (2009). From data to text in the neonatal intensive care Unit: Using NLG technology for decision support and information management. AI Communications, 22(3), 153–186. doi:10.3233/AIC-2009-0453
- Grice, H. P. (1975). Logic and conversation. In Syntax and Semantics 3: Speech Acts (pp. 41–58). Amsterdam: Elsevier.
- Grosz, B. J., Joshi, A. K., & Weinstein, S. (1995). Centering : A Framework for Modeling the Local Coherence of Discourse. Computational Linguistics, 21(2), 203–225.
- Harbusch, K., & Kempen, G. (2009). Generating clausal coordinate ellipsis multilingually: A uniform approach based on postediting. In Proceedings of the 12th European Workshop on Natural Language Generation (ENLG'09) (pp. 138–145). Athens, Greece: Association for Computational Linguistics. Retrieved from http://dl.acm.org/citation.cfm?id=1610219
- Janarthanam, S., & Lemon, O. (2014). Adaptive Generation in Dialogue Systems Using Dynamic User Modeling. Computational Linguistics, 40(4), 883–920. doi:10.1162/COLI_a_00203
- Kibble, R., & Power, R. (2004). Optimizing referential coherence in text generation. Computational Linguistics, 30(4), 401–416. Retrieved from http://www.mitpressjournals.org/doi/abs/10.1162/0891201042544893
- Konstas, I., & Lapata, M. (2013). A global model for concept-to-text generation. Journal of Artificial Intelligence Research, 48, 305–346. doi:10.1613/jair.4025
- Krahmer, E., & Theune, M. (2002). Efficient context-sensitive generation of referring expressions. In K. van Deemter & R. Kibble (Eds.), Information Sharing (pp. 223–264). Stanford, CA: CSLI.
- Krahmer, E., & van Deemter, K. (2012). Computational generation of referring expressions: A survey. Computational Linguistics, 38(1), 173–218. Retrieved from http://www.mitpressjournals.org/doi/abs/10.1162/COLI_a_00088
- Liang, P., Jordan, M. I., & Klein, D. (2009). Learning Semantic Correspondences with Less Supervision. In Proceedings of the Annual Meeting of the Association for Computational Linguistics and International Joint Conference on Natural Language Processing (ACL-IJCNLP'09) (pp. 91–99). Singapore: Association for Computational Linguistics. doi:10.3115/1687878.1687893

References

- Mahamood, S., & Reiter, E. (2011). Generating Affective Natural Language for Parents of Neonatal Infants. In Proceedings of the 13th European Workshop on Natural Language Generation (ENLG'11) (pp. 12–21). Nancy, France: Association for Computational Linguistics. Retrieved from http://www.aclweb.org/anthology/W11-2803
- Mairesse, F., & Walker, M. A. (2011). Controlling User Perceptions of Linguistic Style: Trainable Generation of Personality Traits. Computational Linguistics, 37(3), 455–488.
- McCoy, K. F., McCoy, K. F., Strube, M., & Strube, M. (1999). Generating Anaphoric Expressions: Pronoun or Definite Description? In D. Cristea, N. Ide, & D. Marcu (Eds.), The Relation of Discourse/Dialogue Structure and Reference, Proceedings of the Workshop held in conjunction with ACL-99 (pp. 63–71). Maryland: Association for Computational Linguistics.
- Paiva, D. S., & Evans, R. (2005). Empirically-based control of natural language generation. In Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05) (pp. 58–65). Ann Arbor, Michigan: Association for Computational Linguistics. Retrieved from http://dl.acm.org/citation.cfm?id=1219848
- Searle, J. R. (1969). Speech Acts: An Essay in the Philosophy of Language. Cambridge, UK: Cambridge University Press.
- Walker, M. A., Cahn, J. E., & Whittaker, S. J. (1997). Improvising linguistic style: Social and affective bases for agent personality. In Proceedings of the 1st International Conference on Autonomous Agents (Agents'97) (pp. 96–105). Marina del Rey, CA: ACM Press.
- Walker, M. a., Rainbow, O. C., & Rogati, M. (2002). Training a sentence planner for spoken dialogue using boosting. Computer Speech and Language, 16(3-4), 409–433. doi:10.1016/S0885-2308(02)00027-X