

Einführung in die Pragmatik und Diskurs: Vorlesung 1: Einführung

Vera Demberg, Merel Scholman

Universität des Saarlandes

Sommersemester 2017

Overview for today

- Formalities (and informalities)
- Overview of course topics
- Introduction to discourse structure for language technology

Basic reading:

- Levinson 1983, Chapter 1
- Webber, Egg & Kordoni, 2012.

Woche 1	Intro – VD
Woche 2	Cohesion and Coherence – MS
Woche 3	Frameworks and Mapping – VD
Woche 4	! entfaellt!
Woche 5	Inter-annotator agreement – MS
Woche 6	Discourse processing in humans – MS
Woche 7	Student presentations
Woche 8	Rational Speech Act model – FY
Woche 9	Automatic Relation Classification – VD
Woche 10	Student presentations
Woche 11	! entfaellt!
Woche 12	Student presentations
Woche 13	Discourse representation theory – VD
Woche 14	Student presentations

Introduction to Discourse Structure based on Webber, Egg & Kordoni, 2012

- 1 Where can Discourse Structure help Language Technology?
- 2 History of Computational Discourse modellig

(slides taken from a presentation by Bonnie Webber)

Parsing Text?

To understand where Discourse Structure can help LT, we can start by asking: **Where do discourse and LT come in contact?**

Discourse has long been ignored in training and testing **parsers**:

- That the Penn TreeBank (PTB) files are in reverse chronological order is irrelevant for these tasks,
- As would be scrambling the order in which sentences appear in the files!

Although both discourse and sentence structure can vary with **genre** (eg, *news reports*, *reviews*, *letters*, etc.), parsing a new text benefits more from its **words** having occurred in the training corpus than texts from the same genre [Plank & van Noord, 2011].

Summarizing Text?

Features of discourse structure can contribute to selecting “important” sentences in **text summarization** [Marcu 2000; Schilder 2002; Louis, Joshi & Nenkova 2010].

e.g. Sentences whose content plays the discourse role of

- explanation, or
- comment, or
- example

are considered to be **subordinate**, so may be omitted from extractive summaries [Endres-Niggemeyer, 1998].

Summarizing Text?

- (1) “Mega or non-mega, we feel the prospectus standards need to be considerably improved,” he says. (Implicit = REASON)
“Disclosures are very poor in India.” [wsj_0629]
 - (2) “Disclosures are very poor in India.” (Implicit = INSTANTIATION) He says the big questions – “Do you really need this much money to put up these investments? Have you told investors what is happening in your sector? . . .” – aren't asked of companies coming to market. [wsj_0629]
- ⇒ “Mega or non-mega, we feel the prospectus standards need to be considerably improved,” he says.

Summarizing Text?

Discourse structure can also be used to help repair the anaphoric and coreferential **chaos** of extractive summarization:

- (3) [i] More than 130 bodies are reported to have been recovered after a Gulf Air jet carrying 143 people crashed into the Gulf off Bahrain on Wednesday. [ii] Distraught relatives **also** gathered at Cairo airport, demanding information. [iii] **He also** declared three days of national mourning. [iv] **He** said the jet fell “sharply, like an arrow.” [Otterbacher et al, 2002]

⇒ Unlike parsing, discourse structure **can** potentially benefit summarization.

Analysing and Scoring Student Essays?

Discourse structure is a factor in **assessing the quality** of student essays [Burstein et al 2001; 2003]

Good essays that respond to a **prompt** show clear structure:

- **Introductory material**: segments that provide context for interpreting the thesis, a main idea or the conclusion.
- **Thesis**: segments that state the writers position and are related to the essay prompt.
- **Main idea**: segments that assert the authors main message in conjunction with the thesis.
- **Supporting idea**: segments that support the claims made in the main ideas, thesis statements or conclusions.
- **Conclusion**: segments that summarize the essays argument.

Analysing and Scoring Student Essays?

If such structure is scrambled or difficult to determine, then the quality of an essay suffers.

Sample Prompt [<http://www.ets.org/erater/demo/>]

Often in life we experience a conflict in choosing between something we want to do and something we feel we should do.

In your opinion, are there any circumstances in which it is better for people to do what they want to do rather than what they feel they should do? Support your position with evidence from your own experience or your observations of other people.

Analysing and Scoring Student Essays?

- (4) Throughout our lives, we all find ourselves in a situation at least once, where we have to decide whether we do what we want to or what we feel we should do. This is a very common situation specially among young adults; since we have to decide what we want to make out of our lives. I for instance have to decide to become a lawyer or a doctor. I want to be a lawyer, but I feel I should be a doctor. I can not decide to do what I want or what I think I should since I do not know which is better. [http://www.ets.org/erater/demo/essay_sample_d/]

Thesis (given the prompt)? Main ideas (given the thesis)?

Supporting ideas (given the main ideas)? Conclusion??

⇒ Unlike parsing, discourse structure **can** potentially benefit analysing and scoring student essays.

Information Extraction?

Because within a genre, discourse structure predicts where particular information will be found, if present, it can also potentially benefit **information extraction**.

In descriptions of criminal cases, the victim and perpetrator will be found **before** the alleged offenses and court opinion are detailed [Moens et al 1999; 2000].

In a letter, the writer's name comes at the **end** of the text.

Opinion Mining and Sentiment Detection?

Because within a genre, discourse structure predicts how information should be interpreted, if present it can also potentially benefit **opinion mining and sentiment detection**.

e.g, Evaluation expressions

- at the **end** of a review (ie, summarizing the writer's opinion)
- in a **prominent position** (eg, paragraph-initial)

are generally better predictors of a writer's overall opinion than those found elsewhere [Taboada et al 2009].

Statistical Machine Translation?

Discourse structure has long been ignored in **SMT**.

But preliminary evidence shows that improvements in translating **anaphoric expressions** (with neither phrase-local or tree-local antecedents) can improve SMT [Le Nagard & Koehn 2010; Hardmeier & Federico 2010].

Because these improvements do not lead to improved Bleu scores other metrics are needed in order to assess them [Hardmeier & Federico 2010].

⇒ Even in **SMT**, discourse structure can deliver potential benefit.

Early Computational Discourse Modelling

If we agree that taking account of discourse structure can help LT, why has it not yet really done so?

Start by looking at the **history** of computational work on discourse modelling.

Early computational work generally assumed that discourse had an underlying **tree structure**, similar to the parse tree of a sentence.

At issue was what the internal nodes of the tree and its other formal properties corresponded to.

Rhetorical Structure Theory (RST)

Rhetorical Structure Theory [Mann & Thomson, 1988] associates a discourse with a tree structure through Context-Free (CF) rewrite rules called **schemas**.


An RST analysis covers the discourse with a tree structure, much as a syntactic parse tree covers a sentence.

Dominance of a node over its children corresponds to a rhetorical relation holding between the text units associated with those child nodes (which project to adjacent text spans).


Precedence between nodes corresponds to their order in the text.

RST Schemas as CF rules

RST notation





Standard tree notation



Example: RST Analysis

- (5) a. You should come visit.
b. Edinburgh is lovely in early fall
c. and there are no rabbits around.



EPICURE [Dale, 1992]

In [Dale, 1992], complex task instructions derive from


- 1 a tree structure produced through recursive CF-decomposition of a top-level task into sub-tasks, with **precedence** corresponding to temporal order and **dominance**, to sub-task inclusion;
- 2 aggregation of similar sisters for realization in a single clause.

Each internal node in the tree corresponded to the next step in a plan to accomplish its parent.

EPIASURE [Dale, 1992]

[Make bean soup]	→	[Prep ingredients]	[Cook ingredients]		
[Prep ingredients]	→	[Prep beans]	[Prep veg]		
[Prep beans]	→	Soak beans	Drain beans	Rinse beans	
[Prep veg]	→	[Prep onion]	[Prep potato]	[Prep carrots]	
		Slice celery			
[Prep onion]	→	Peel onion	Chop onion		
[Prep potato]	→	Peel potato	Chop potato		
[Prep carrots]	→	Scrape carrots	Chop carrots		
[Cook ingredients]	→	[Stage 1]	[Stage 2]		
[Stage 1]	→	Melt butter	Add veg	Saute veg	
		Add beans, stock, milk	Simmer		
[Stage 2]	→	Liquidise	Stir in cream	Season	Reheat

EPICURE [Dale, 1992]



Generation via top-down LR traversal [Dale, 1992]

Soak, drain and rinse the butter beans. ($e_7 - e_9$)

Peel and chop the onion. ($e_{23} - e_{24}$)

Peel and chop the potato. ($e_{25} - e_{26}$)

Scrape and chop the carrots. ($e_{27} - e_{28}$)

Slice the celery. (e_{13})

Melt the butter. (e_{14})

Add the vegetables. (e_{15})

Saute them. (e_{16})

Add the butter beans, the stock and the milk. (e_{17})

Simmer. (e_{18})

Liquidise the soup. (e_{19})

Stir in the cream. (e_{20})

Add the seasonings. (e_{21})

Reheat. (e_{22})

Intentional Discourse Structure [Grosz and Sidner, 1986]

Grosz & Sidner (1986) posit a tree structure for the *intentional structure* of discourse:

- Nodes correspond to **speaker intentions**.
- **Dominance** in the tree corresponds to the intention of a daughter node supporting that of its parent;
- **Precedence** corresponds to the need to satisfy an earlier intention before one that follows.

* * * * *

Lost in this early work on the tree structure of discourse was a **linear** model [Sibun, 1992], that seemed to provide a simpler account of certain types of expository text.

Current Computational Discourse Modelling

What has happened since then?

- Continued work on mainly tree-like discourse structures [Asher & Lascarides, 2003; Polanyi et al, 2004] and on some more complex graph structures [Wolf & Gibson, 2005]
- Work on **Topic structure** of discourse
- Work on **Functional structure** of discourse
- Work on **Coherence relations** – “higher-order” pred-arg structure of discourse.

This discourse modelling tends to be data-intensive, reflecting the goal of robust applications.

Topic Structure

Expository text can be viewed as a sequence of **topically coherent** segments, whose order may become conventionalized over time:

	Wisconsin	Louisiana	Vermont
1	Etymology	Etymology	Geography
2	History	Geography	History
3	Geography	History	Demographics
4	Demographics	Demographics	Economy
5	Law and government	Economy	Transportation
6	Economy	Law and government	Media
7	Municipalities	Education	Utilities
8	Education	Sports	Law and government
9	Culture	Culture	Public Health
10

Wikipedia articles about US states

Topic Structure

Being able to recognize topic structure was originally seen as benefitting **information retrieval** [Hearst, 1994; 1997].

Recent interest comes from the potential use of topic structure in **segmenting lectures, meetings or other speech events**, making them more amenable to search [Galley 2003; Malioutov & Barzilay 2006].

Topic Structure

Computational approaches to topic structure and segmentation assume that:

- The topic of each discourse segment relates to the topic of the discourse as a whole (eg, History of Vermont → Vermont).
- The only relation holding between sister segments, if any, is pure sequence, although certain sequences may be more common than others (cf. Wikipedia articles).
- The topic of a segment differs from those of its adjacent sisters. (Adjacent spans that share a topic are taken to belong to the same segment.)
- Topic predicts lexical choice, either of all words of a segment or just its content words (ie, excluding “stop-words”).

Topic Structure

- Making this structure explicit (ie, topic segmentation) uses either
- **semantic-relatedness**, where each segment is taken to consist of words that relate to each other more than to words outside the segment [Hearst 1994, 1997; Choi et al 2001; Bestgen 2006; Galley et al 2003; Malioutov & Barzilay 2006]
 - **topic models**, where each segment is taken to be produced by a distinct, compact lexical distribution [Purver et al, 2006; Eisenstein & Barzilay 2008; Chen et al 2009].

[Purver, 2011] contains a useful overview and survey of this work.

Topic Segmentation through Semantic-relatedness

All computational models that use semantic-relatedness for topic segmentation have:


- 1 a **metric** for assessing the semantic relatedness of terms within a proposed segment;
- 2 a **locality** that specifies what units within the text are assessed for semantic relatedness;
- 3 a **threshold** for deciding how low relatedness can drop before it signals a shift to another topic.

Text Tiling [Hearst 1994, 1997]

- 1 **Metric:** *Cosine similarity*, using a vector representation of fixed-length spans (pseudo-sentences) in terms of frequency of word stems (ie, words from which inflection has been removed)
- 2 **Locality:** Cosine similarity is computed between adjacent spans (and only adjacent spans)
- 3 **Threshold:** Empirically-determined in order to select where to place segment boundaries.

TextTiling – Computed similarity of adjacent blocks

TextTiling a popular science article. Vertical lines show manually-assigned topic boundaries. Peaks indicate coherency, and valleys, potential breaks between tiles. [http://people.ischool.berkeley.edu/~hearth/papers/subtopics-sigir93/sigir93.html]



Topic Segmentation through Topic Models

Topic segmentation using topic models can take advantage of both

- Features internal to a segment (*Segmental Features*), including words (all words or just content words) and syntax
- Features occurring at segmental boundaries (*Boundary Features*), including **discourse cue words** (eg, “now”, “so”, “anyway”), syntax and (in speech) pauses and intonation.

N.B. What cue words are indicating might be better described as **functional structure** than **topic structure**.

Functional structure

Texts within a given genre – eg,

- news reports
- errata
- scientific papers
- letters to the editor of the *New York Times*
- ...

generally share a similar structure, that is independent of topic (eg, sports, politics, disasters; or molecular genetics, radio astronomy, SMT), instead reflecting the **function** played by their parts.

Example: News Reports

Best known is the **inverted pyramid** structure of news reports:

- Headline
- Lead paragraph, conveying **who** is involved, **what** happened, **when** it happened, **where** it happened, **why** it happened, and (optionally) **how** it happened
- Body, providing more detail about who, what, when, . . .
- Tail, containing less important information

This is why the first (ie, lead) paragraph is usually the best *extractive summary* of a news report.

Example: Errata

Also recognizable are **errata** – declarations of errors made in previous issue of a periodical and correct versions:

- **Correct statement**
 - *Description of error*
- (6) **EMPIRE PENCIL, later called Empire-Berol, developed the plastic pencil in 1973.** *Yesterday's Centennial Journal misstated the company's name. [wsj_1751]*
- (7) **PRINCE HENRI is the crown prince and hereditary grand duke of Luxembourg.** *An article in the World Business Report of Sept. 22 editions incorrectly referred to his father, Grand Duke Jean, as the crown prince. [wsj_1871]*

Example: Scientific articles/abstracts

Well-known in academia is the multi-part structure of scientific papers (and, more recently, their abstracts):

- **Objective** (aka *Introduction, Background, Aim, Hypothesis*)
- **Methods** (aka *Method, Study Design, Methodology, etc.*)
- **Results** or *Outcomes*
- **Discussion**
- Optionally, **Conclusions**

N.B. Not every sentence within a section need realise the same function: Fine-grained functional characterizations of scientific papers show them serving a range of functions [Liakata, 2010].

Functional Structure

Automatic annotation of functional structure is seen as benefitting:

- Information extraction: Certain types of information are likely to be found in certain sections [Moens 1999; 2000].
- Extractive summarization: More “important” sentences are more likely to be found in certain sections.
- Sentiment analysis: Words that have an objective sense in one section may have a subjective sense in another [Taboada, 2009].
- Citation analysis: A citation may serve different functions in different sections [Teufel, 2010].

Functional structure

Computational approaches to functional structure and segmentation assume that:

- The function of a segment relates to that of the discourse as a whole.
- While relations may hold between sisters (eg, *Methods* constrain *Results*), only sequence has been used in modelling.
- Function predicts more than lexical choice:
 - indicative phrases such as “results show” (\rightarrow *Results*)
 - indicative stop-words such as “then” (\rightarrow *Method*).
- Functional segments usually appear in a specific order, so either sentence position is a feature used in modelling or sequential models are used.

Functional structure

The internal structure of segments has usually been ignored in high-level functional segmentation [Chung, 2009; Lin et al, 2006; McKnight, 2003; Ruch, 2007].

But given the results of work in fine-grained modelling of functional structure, it is not surprising that Hirohata et al (2008) found significant boundary features:

- Properties of the first sentence of a segment differ from those of the rest (as in 'BIO' approaches to Named Entity Recognition).
- Modelling this leads to improved performance in high-level functional segmentation (ie, 94.3% per sentence accuracy vs. 93.3%).

Labelled biomedical abstracts

Much function-based modelling has been on biomedical text [Chung, 2009; Guo et al, 2010; Hirohata et al, 2008; Liakata et al, 2010; Lin et al, 2006; Mcknight & Srinivasan, 2003; Ruch et al, 2007], where texts with explicitly labelled sections serve as **free training data** for segmenting unlabelled texts.

- (8) **BACKGROUND:** Mutation impact extraction is a hitherto unaccomplished task in state of the art mutation extraction systems. ... **RESULTS:** We present the first rule-based approach for the extraction of mutation impacts on protein properties, categorizing their directionality as positive, negative or neutral. ... **CONCLUSION:** ... Our approaches show state of the art levels of precision and recall for Mutation Grounding and respectable level of precision but lower recall for the task of Mutant-Impact relation extraction. ... [PMID 21143808]

Unlabelled biomedical abstracts

- (9) We propose two methods for finding similarities in protein structure databases. Our techniques extract feature vectors on triplets of SSEs (Secondary Structure Elements) of proteins. These feature vectors are then indexed using a multidimensional index structure. Our first technique considers the problem of finding proteins similar to a given query protein in a protein dataset. This technique quickly finds promising proteins using the index structure. These proteins are then aligned to the query protein using a popular pairwise alignment tool such as VAST. We also develop a novel statistical model to estimate the goodness of a match using the SSEs. Our second technique considers the problem of joining two protein datasets to find an all-to-all similarity. Experimental results show that our techniques improve the pruning time of VAST 3 to 3.5 times while keeping the sensitivity similar. [PMID 16452789]

Other work on functional segmentation of legal texts [Moens 1999, 2000] and student essays [Burstein et al 2001, 2003].

“Higher-order” pred-arg structures

Discourse also has structure arising from semantic and pragmatic relations that hold between the referents of its clauses.

These resemble relations between the referents of NPs that serve as **args** to a **predicate** conveyed by a verb (**PropBank**) or noun (**NomBank**).

These “higher-order” pred-arg structures (**discourse relations** or **coherence relations**) are usually signalled by a discourse connective

- a conjunction like *because* or *but*,
- a discourse adverbial like *nevertheless* or *instead*.

Discourse Connectives \neq Discourse Cues

Words like English *so* are ambiguous, with some tokens that serve as discourse connectives and some, as discourse cues.

(10) But C.J.B. Marshall, vicar of a nearby church, feels the fault is in the stairs from the bell tower that are located next to the altar. “**So** crunch, crunch, crunch, bang, bang, bang – here come the ringers from above, making a very obvious exit while the congregation is at prayer,” he says. [wsj_0089]

(11) Indeed, Judge O'Brien ruled that “it would be easy to conclude that the USIA’s position is ‘inappropriate or even stupid,’ ” but it’s the law. So the next step, I suspect, is to try to get the law changed. [wsj_0108]

N.B. Discourse relations may also be signalled in other ways, like English *that means*, *what if*, etc. [Prasad et al, 2008]).

“Higher-order” pred-arg structures

But just as relations between referents can be signalled purely by adjacency — cf. **noun-noun modifier** constructions in English

- human-computer communication
- container ship crane operator courses

coherence relations can be conveyed through adjacency between clauses or sentences (aka *implicit connectives*).

(12) Viewers may not be cheering, either.

Soaring rights fees will lead to an even greater clutter of commercials. [wsj_1057]

“Higher-order” pred-arg structures

But just as relations between referents can be signalled purely by adjacency — cf. **noun-noun modifier** constructions in English

- human-computer communication
- container ship crane operator courses

coherence relations can be conveyed through adjacency between clauses or sentences (aka *implicit connectives*).

(13) Viewers may not be cheering, either. (implicit=REASON)
Soaring rights fees will lead to an even greater clutter of commercials. [wsj_1057]

Resources annotated with Coherence Relations

The Penn Discourse TreeBank is currently the largest resource manually annotated for discourse connectives, their arguments, and the senses they convey [Prasad, 2008].

PDTB Relations	No. of tokens
Explicit	18459
Implicit	16224
...	...

Resources annotated with Coherence Relations

Resources annotated like the PDTB are being created for:

- Modern Standard Arabic [Al-Saif and Markert, 2010]
- Chinese [Xue, 2005, 2010]
- Czech [Mladová et al, 2008]
- Danish and Italian [Buch-Kromann and Korzen, 2010]
- Dutch [van der Vliet et al, 2011]
- German [Stede, 2004, 2008]
- Hindi [Oza et al, 2009]
- Turkish [Zeyrek et al, 2010]

In all these language, coherence relations have two (and only two) arguments.