Abduce or not Abduce ? That is the Question.
On the Role of Abductive Reasoning within Discourse Interpretation.

Matthias Hartung
University of Heidelberg
Department of Computational Linguistics
hartung@cl.uni-heidelberg.de

Philipp Cimiano
University of Karlsruhe
Institute AIFB
cimiano@aifb.uni-karlsruhe.de

April 13, 2007

Outline

Discourse Interpretation: Integration vs. Separation Hypothesis

Discourse Interpretation via Coherent Discourse Structure

- common ground between Hobbs (1993) and Asher&Lascarides (2003): discourse interpretation via a coherent discourse structure which is modelled by means of discourse relations connecting single discourse segments
- causes of debate:
  - How should the computation of such a coherent discourse structure be carried out ?
  - Which role in this task should be played by abductive reasoning ?

Separation Hypothesis – Asher & Lascarides (2003)

- modularized perspective: computation of *discourse structure* and *discourse interpretation* are sketched as strictly separated tasks !
- coherent discourse structure is modelled by discourse relations connecting discourse segments
- semantic effects of discourse relations: enriching the compositional semantics of the sentences within a discourse by additional implicit meaning

Integration Hypothesis – Hobbs (1993)

- 'To interpret a text, one must prove the logical form of the text from what is already mutually known, allowing for coercions, merging redundancies where possible, and **making assumptions where necessary**'. (Hobbs 1993)

- Integration of additional assumptions yields valid proofs of the following form within an abductive calculus:

$$\frac{\begin{array}{c}\forall x : p(x) \supset q(x) \\ \exists q(A)\end{array}}{\exists p(A)}$$

- problem: *linguistic* constraints on weighted abduction ??

Motivating Data

Experiment

- Example (cp. Danlos 2001):
  π_1 : *Fred sprang aus einem Flugzeug ab.* π_2 : *Er starb.*
- Questions:
  (Q1) Was war die Todesursache ?
  (Q2) Flog das Flugzeug zum Zeitpunkt des Absprungs ?
  (Q3) Benutzte Fred einen Fallschirm ?
- Results:
  - `Q1` as control variable: 87.5% of answers related to jump injuries $\Rightarrow$ *causal interpretation* of the connection between π_1 and π_2 !
  - *among those*:
    `Q2 = '(probably) yes'`: 92.9%
    `Q3 = '(probably) no'`: 64.2%

Conclusion

- *„Wenn beide Sachverhalte nichts miteinander zu tun haben, dann ist genauso wahrscheinlich, dass er einen Fallschirm getragen hat, wie dass er keinen getragen hat. Wenn er aufgrund des Sprungs gestorben ist, ist es wahrscheinich, dass er keinen Fallschirm getragen hat. Es kann aber auch sein, dass er nicht in der Lage war, damit umzugehen und doch einen benutzt hat.“*
- Hypothesis: Human interpreters are capable of interpreting discourse in such a way that the semantic consequences of discourse relations override default (a priori) world knowledge. Thus, the evaluation of discourse relations enables us to extract implicit information from texts.
- Question: **What kind of knowledge is required to infer discourse relations ??**

SDRT's Approach to Discourse Interpretation

Typology of Discourse Relations in SDRT – Glue Logic (1)

RESULT/EXPLANATION:

Trigger:

- $(^{?}R(\alpha, \beta) \wedge cause_D(\alpha, \beta)) > \textsc{Result}(\alpha, \beta)$
- $(^{?}R(\alpha, \beta) \wedge cause_D(\beta, \alpha)) > \textsc{Explanation}(\alpha, \beta)$

Semantic Effects:

- $\textsc{Explanation}(\alpha, \beta) \Rightarrow (\neg e_\alpha \prec e_\beta)$
- $\textsc{Explanation}(\alpha, \beta) \Rightarrow (event(e_\beta) \Rightarrow e_\beta \prec e_\alpha)$

Typology of Discourse Relations in SDRT – Glue Logic (2)

ELABORATION:

Trigger:

- $(^{?}R(\alpha, \beta) \wedge subtype(\beta, \alpha)) > \text{ELABORATION}(\alpha, \beta)$
- $(^{?}R(\alpha, \beta) \wedge coref(\beta, \alpha)) > \text{ELABORATION}(\alpha, \beta)$

Semantic Effects:

- $\text{ELABORATION}(\alpha, \beta) \Rightarrow \textit{Part-of}(e_\beta, e_\alpha)$

Typology of Discourse Relations in SDRT – Glue Logic (3)

NARRATION:

Trigger:

- $({}^{?}R(\alpha,\beta) \wedge occasion(\alpha,\beta)) > \textsc{Narration}(\alpha,\beta)$
- $({}^{?}R(\alpha,\beta) \wedge [\textit{fall}(e_1,x))](\alpha) \wedge [\textit{help-up}(e_2,x,y))](\beta) > occasion(\alpha,\beta))$

Semantic Effects:

- $\textsc{Narration}(\alpha,\beta) \Rightarrow overlap(prestate(e_\beta), poststate(e_\alpha))$
- $\textsc{Narration}(\alpha,\beta) \Rightarrow K_\alpha \sqcap K_\beta$ (Topic Constraint)

Knowledge Sources

'Some of the axioms of the Glue Logic appeal to certain semantic relations as clues for rhetorical relations.' (Asher & Lascarides 2003)

SDRT makes use of the following knowledge sources in order to infer these semantic relations:

- lexically specified knowledge
- defeasible inferences from lexical information
- subcategorization frames
- subtype information
- world knowledge

Example: Lexically Specified Knowledge

Example: Mary was annoyed. Peter didn't call.

$$\lambda e \lambda x \lambda y \quad \begin{array}{|l|} \hline \pi_1, \pi_2 \\ \hline \pi_1\text{:} \begin{array}{|c|} \hline \\ \hline \textit{annoy(e,u,y)} \\ \hline \end{array} \quad \pi_2\text{:} \begin{array}{|c|} \hline \\ \hline \textit{experience}(^{?}\textit{(u),y)} \\ \hline \end{array} \\ \textsc{Explanation}(\pi_1, \pi_2) \\ \hline \end{array}$$

$\Rightarrow$ no prototypical semantic knowledge about the relation between *annoy* and *call* involved; instead: any event e' that y experiences as an undergoer will do !

Example: Defeasible Inferences from Lexical Information

Example: Max fell. John hit him.

$$\lambda e \lambda x \lambda y \; \begin{array}{|l|} \hline \pi_1, \pi_2 \\ \hline \pi_1: \begin{array}{|l|} \hline \\ \hline \textit{hit}(e, x, y) \\ \hline \end{array} \quad \pi_2: \begin{array}{|l|} \hline e' \\ \hline ?^{Move}(e', y) \\ \hline \end{array} \\ \textit{Def-Causal-Conseq}(\pi_1, \pi_2) \\ \hline \end{array}$$

$\Rightarrow$ *hitting* normally entails a *movement* of the person or artifact being hit; combination of underspecified lexical information and semantic type hierarchy leads to the solution !

The Need for Abduction

The Need for Abduction

Examples (cp. Danlos 2001):

- π_1 : John jumped off a plane. π_2 : He died.
- π_3 : John forgot his parachute. π_4 : He died.

$\Rightarrow$ Neither lexical information nor prototypical world knowledge suggests a causal relation between π_1 and π_2 or π_3 and π_4, respectively !

$\Rightarrow$ Contradicts our empirical findings concerning human interpretation capabilities !

$\Rightarrow$ Our solution: Process linguistic and world knowledge within an abductive calculus in order to infer discourse relations via defeasible inferences where necessary !

Our Approach: Integrating SDRT and Abduction in an Ontology-based Account

Ingredients

- ontology containing linguistic knowledge about semantics of events
  - temporal structure of sub-events (aktionsarten; cp. Vendler 1967, Moens & Steedman 1988)
  - relations between sub-events and event's roles (cp. Pustejovsky 1995)
  - contextual compositionality of events (cp. Moens & Steedman 1988)
- world knowledge
- DRT as knowledge representation formalism
- underlying proof system: first-order calculus operating on DRSs (Kamp & Reyle 1996), extended by abductive techniques

Definition of Concepts along their Pre- and Post-Conditions

- informally: 'Events describe processes in the course of which their participants undergo a change of states.'
- **Pre- and Post-Conditions of events (cp. Allen 1994):**
  s_1: state of participants at t_1, i.e. the time point immediately before the event takes place
  s_2: state of participants at t_2, i.e. the time point immediately after the event has taken place
- temporal ordering: $s_1 \supset\subset e \supset\subset s_2$
- notion of $\supset\subset$ ('abut'): $\forall e_1, e_2$:
  $(e_1 \supset\subset e_2) \iff (e_1 \prec e_2) \wedge \neg\exists e_3 : (e_1 \prec e_3 \prec e_2)$

DRT as Representation Formalism – Example

$$K_O: \begin{array}{|l|}\hline e,x \\ \hline \textit{ScoreGoal(e,x)} \\ \\ \hline \end{array} \Rightarrow K_O^{Pre}: \begin{array}{|l|}\hline s_1, b, g \\ \hline \textit{BallPossession(x,b)} \\ \neg\textit{BallPosition(b,g)} \\ \textit{Ball(b)} \\ \textit{Goal(g)} \\ s_1 \supset\subset e \\ \\ \hline \end{array}$$

$$K_O: \begin{array}{|l|}\hline e,x \\ \hline \textit{ScoreGoal(e,x)} \\ \\ \hline \end{array} \Rightarrow K_O^{Post}: \begin{array}{|l|}\hline s_2, b, g \\ \hline \textit{BallPosition(b,g)} \\ \neg\textit{BallPossession(x,b)} \\ \textit{Ball(b)} \\ \textit{Goal(g)} \\ e \supset\subset s_2 \\ \\ \hline \end{array}$$

$$K_O: \begin{array}{|l|}\hline e,x \\ \hline \textit{ScoreGoal(e,x)} \\ \\ \hline \end{array} \Rightarrow K_O^{Sub}: \begin{array}{|l|}\hline e' \\ \hline \textit{Shot(e',x)} \\ \\ \hline \end{array}$$

Inferring Discourse Relations from the Pre- and Post-Conditions of Events

In our approach, the presence of a discourse relation R between two events e_1 and e_2 is defined in terms of the specific configuration of the pre- and post-conditions of e_1 and e_2.

- $(K_1^{Post} \Rightarrow_{GMP} K_2^{Pre}) \Rightarrow occasion(K_1, K_2)$
- $((K_1^{Post} \Rightarrow_{GMP} K_2^{Post}) \wedge ArgumentCoherence(e_1, e_2)) \Rightarrow cause_D(K_1, K_2)$
- $((K_1^{Pre} \Rightarrow_{GMP} K_2^{Pre}) \wedge (K_1^{Post} \Rightarrow K_2^{Post})) \Rightarrow coref(K_1, K_2)$
- $(K_2 \Rightarrow_O K_1^{Sub}) \Rightarrow subtype(K_1, K_2)$

Example in More Detail (1)

π_1 : Fred jumped off a plane. π_2 : He died.

$$K_{\pi_1} : \begin{array}{|l|}\hline \textit{e,x,y,y'} \\ \hline \textit{jump-off(e,x,y)} \\ \\ \hline \end{array} \Rightarrow K^{Pre}_{\pi_1} : \begin{array}{|l|}\hline s_1 \\ \hline \textit{be-at(x,y)} \\ \textit{highAltitude(y)} \\ s_1 \supset\subset e \\ \\ \hline \end{array} \wedge K^{Post}_{\pi_1} : \begin{array}{|l|}\hline s_2 \\ \hline \textit{be-at(x,y')} \\ \textit{lowAltitude(y')} \\ y \neq y' \\ e \supset\subset s_2 \\ \\ \hline \end{array}$$

$$K_{\pi_2} : \begin{array}{|l|}\hline e, x \\ \hline \textit{die}(e, x) \\ \\ \hline \end{array} \Rightarrow K^{Pre}_{\pi_2} : \begin{array}{|l|}\hline s_1 \\ \hline \neg\textit{dead(x)} \\ s_1 \supset\subset e \\ \\ \hline \end{array} \wedge K^{Post}_{\pi_2} : \begin{array}{|l|}\hline s_2 \\ \hline \textit{dead(x)} \\ e \supset\subset s_2 \\ \\ \hline \end{array}$$

Outline
Discourse Interpretation: Integration vs. Separation Hypothesis
Motivating Data
SDRT's Approach to Discourse Interpretation
The Need for Abduction
Our Approach: Integrating SDRT and Abduction in an Ontology-based Account
Conclusion

Example in More Detail (2) – 'The Magic Part'

π_1 : Fred jumped off a plane. π_2 : He died.

$$K_{\pi_1} : \begin{array}{|l|}\hline \textit{e,x,y,y'} \\ \hline \textit{jump-off(e,x,y)} \\ \\ \hline \end{array} \Rightarrow K^{Pre}_{\pi_1} : \begin{array}{|l|}\hline s_1 \\ \hline \textit{be-at(x,y)} \\ \textit{highAltitude(y)} \\ s_1 \supset\subset e \\ \\ \hline \end{array} \wedge K^{Post}_{\pi_1} : \begin{array}{|l|}\hline s_2 \\ \hline \textit{be-at(x,y')} \\ \textit{lowAltitude(y')} \\ \textbf{\textit{dead(x)}} \\ y \neq y' \\ e \supset\subset s_2 \\ \\ \hline \end{array}$$

$$K_{\pi_2} : \begin{array}{|l|}\hline e, x \\ \hline \textit{die}(e, x) \\ \\ \hline \end{array} \Rightarrow K^{Pre}_{\pi_2} : \begin{array}{|l|}\hline s_1 \\ \hline \neg\textit{dead(x)} \\ s_1 \supset\subset e \\ \\ \hline \end{array} \wedge K^{Post}_{\pi_2} : \begin{array}{|l|}\hline s_2 \\ \hline \textbf{\textit{dead(x)}} \\ e \supset\subset s_2 \\ \\ \hline \end{array}$$

- Try to accomodate additional premises until $K^{Post}_{\pi_1} \Rightarrow_m K^{Post}_{\pi_2}$!
- Plausibility check: *Backward Chaining* of axioms representing world knowledge

Plausibility Checks by Backward Chaining

1. Possible reasons for being dead ?
   $jump-off(x,y) \wedge highAltitude(y) \supset dead(x)$

2. Possible verification of $highAltitude(y)$?
   $plane(y) \wedge fly(y) \supset highAltitude(y)$

$\Rightarrow$ $fly(p)$ has to be *assumed* instead of being verified by the context !
$\Rightarrow$ implicit knowledge human interpreters are capable to infer !

Further Work

Is there a second layer of constraints on the 'magic part' coming from *linguistic knowledge* ??

$$K_{O_1} : \begin{array}{|l|}\hline \textit{e,x,y,y'} \\ \hline \textit{jump-off(e,x,y)} \\ \\ \hline\end{array} \Rightarrow K^{Pre}_{O_1} : \begin{array}{|l|}\hline s_1 \\ \hline \textit{be-at(x,y)} \\ s_1 \supset\subset e \\ \\ \hline\end{array} \wedge K^{Post}_{O_1} : \begin{array}{|l|}\hline s_2 \\ \hline \textit{be-at(x,y')} \\ {}^{?}\boldsymbol{\textit{state(U:x)}} \\ y \neq y' \\ e \supset\subset s_2 \\ \\ \hline\end{array}$$

Conclusion

To cut a long story short...

- Empirical findings suggest a divergence between human interpreters' capabilities and the modelling power of SDRT.
- Viable solution: extension of an ontology-based approach to discourse structure by abductive reasoning
- Suggestion: Abduction is – as opposed to artifical weights (cp. Hobbs 1993) – constrained by world knowledge and linguistic knowledge.

Acknowledgement

Part of this work has been carried out during the first author's occupation as a research assistant at Institute AIFB. The financial support granted by AIFB is gratefully acknowledged.