CENTRE FOR LANGUAGE TECHNOLOGY
MACQUARIE UNIVERSITY – SYDNEY

The Semantic Representation of Temporal Expressions in Text

Robert Dale and Pawel Mazur
rdale | mpawel@ics.mq.edu.au

Agenda

- Background
- The DANTE System and Its Architecture
- Local Semantics
- Evaluation
- Error Analysis
- Future Work

Background

- References to time – temporal expressions – are widespread in text
- For many tasks in involving text processing we need to know when events take place

→ Two key tasks:

  – The recognition and interpretation of temporal expressions

  – Time-stamping of events described in text

Precise Temporal Expressions #1: Points in Time

- Fully specified dates and times
  - 21st November 1994; 5:16am on November 21st 1994; the 1990s
- Underspecified dates and times:
  - January 3; 9 pm; ten minutes to 3; the '60s
- Relative dates and times:
  - today; last month; sixty seconds later; at 6 a.m. today

Precise Temporal Expressions #2: Durations, Sets, Frequencies

- Durations:
  - three-hour; the past four years; the next two weeks
- Sets of points in time and frequencies:
  - every Tuesday in 1999; monthly; the first three days of every month; the last two Fridays of every second month

Other Temporal Expressions #1

- Fuzzy Specifications:
  - past, present, future, seasons, fiscal years, year quarters and halves, weekends, mornings, afternoons …
- Modified Expressions:
  - more than a decade ago; less than 2 hours long
- Non-Specific Expressions:
  - Winters are cold here.
  - I love December.

Other Temporal Expressions #2

- Event-Anchored Specifications:
  - I remember the day that Roosevelt died.
  - The firefighters came home three days after the fire.
- Culturally-Determined Specifications:
  - She was not allowed to play video games last school year, and her grades improved dramatically.

The ACE TERN Task

- TERN = Time Expression Recognition and Normalisation
- Given a temporal expression in a document, annotate this with an ISO date-time object (TIMEX2)
  - The flight arrives at <TIMEX VAL="2007-12-06T10:00AEST"> 10 in the morning on 6th December, Sydney time</TIMEX>.
  - The conference will be over <TIMEX VAL="2007-12-07"> tomorrow</TIMEX>.
  - The trip lasted <TIMEX VAL="P2W">2 weeks</TIMEX>.
  - He hasn't been here in <TIMEX VAL="P2.5M" ANCHOR_VAL:="2000-10-25" ANCHOR_DIR="ENDING">two and a half months</TIMEX>

Agenda

- Background
- The DANTE System and Its Architecture
- Local Semantics
- Evaluation
- Error Analysis
- Future Work

The DANTE System

- DANTE = Detection And Normalisation of Time Expressions
- Modular design: Recognizer, Interpreter and Normaliser
- Implemented in Java, using GATE libraries
- Document processing organised as a pipeline
- A Rule-based system, with a recognition grammar implemented in JAPE: 250 rules, 80 macros, 31 gazetteers with a total of 1418 entries
- An intermediate representation of local semantics used as the interface between detection and interpretation

Steps in Time Expression Recognition and Normalisation

- Temporal Expression Recognition:
  - detect the extent of the temporal expression in the text and provide its local semantics
- Temporal Expression Interpretation:
  - use information from the document context to turn the recognized expression into a fully specified date and time
- Temporal Expression Normalisation:
  - normalise this fully specified date and time to a predefined time zone

DANTE's Architecture

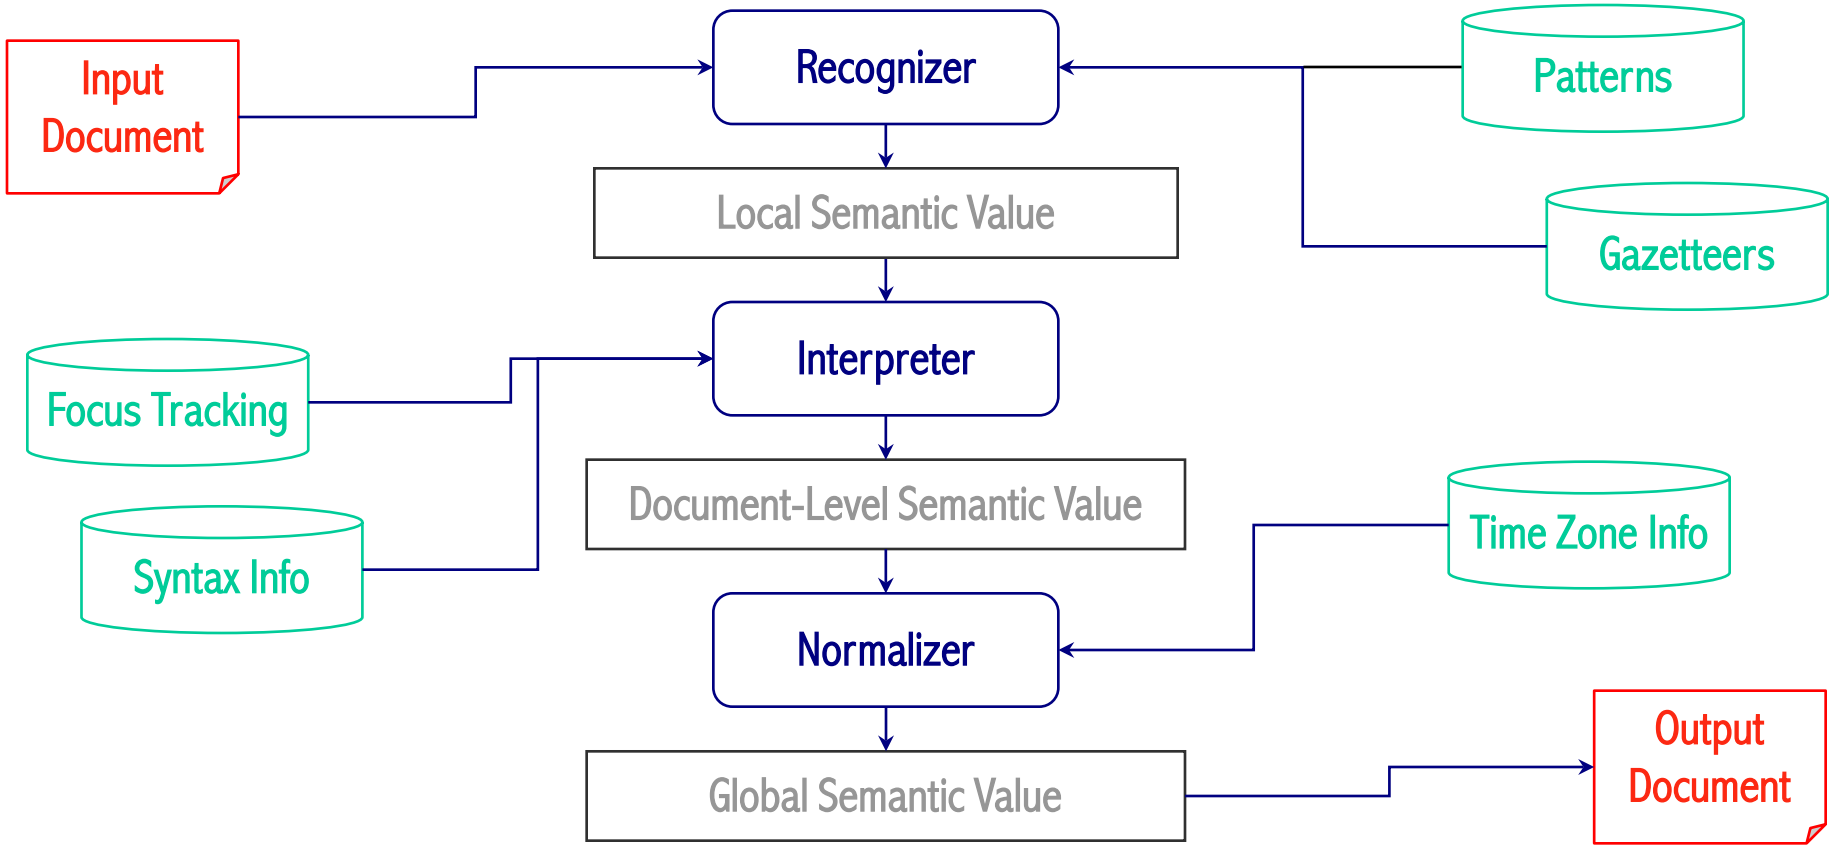

The Steps in DANTE's Processing Pipeline

1. Tokenisation
2. Gazetteer Lookup
3. Sentence Splitter
4. POS tagger (Hepple)
5. Time Expression Recognizer
6. Dependency Parser (e.g. Minipar or Connexor)*
7. Parser-based Time Expression Recognizer*
8. Time Expression Interpreter
9. Time Expression Normaliser*

* optional, not used in ACE evaluations

Agenda

- Background
- The DANTE System and Its Architecture
- Local Semantics
- Evaluation
- Error Analysis
- Future Work

Local Semantic Interpretation

- Observation: a principled distinction can be made between
  - the interpretation of the semantics of a temporal expression devoid of its context of use, and
  - the fuller interpretation of that expression when the context is taken into account
- Benefits:
  - Separates different kinds of processing
  - Allows separate testing and evaluation of distinct components
  - Increases modularity of software and possibility of interchange

Local Semantic Interpretation

- Temporal expressions have a context-independent interpretation:
  - yesterday always means the day before today
  - Thursday is always the fourth day in a week
  - July is always the seventh month in a year
  - next year always means the year following the reference year
- Local semantics:
  - derivable in a compositional manner from the components of the expression
- Global semantics:
  - may require arbitrary inference, reasoning and calendar arithmetic

Representing Temporal Semantics

$$
\begin{bmatrix}
\textit{point} & & & \\
\text{TIMEANDDATE} & \text{TIME} & \text{HOUR} & \textit{15} \\
 & & \text{MINS} & \textit{00} \\
 & & \text{SECS} & \textit{00} \\
 & \text{DATE} & \text{DAY} & \text{DAYNAME} \quad \textit{D4} \\
 & & & \text{DAYNUM} \quad \textit{13} \\
 & & \text{MONTH} & \textit{5} \\
 & & \text{YEAR} & \textit{2006} \\
\text{ZONE} & \textit{Z} & &
\end{bmatrix}
$$

Underspecification

- **We will meet on Thursday.**

- **We will meet on 13th May.**

Compositional Semantics by Unification

- **We will meet on Thursday 13th May.**

$$\begin{bmatrix} \textit{point} \\ \text{TIMEANDDATE} & \begin{bmatrix} \text{DATE} & \begin{bmatrix} \text{DAY} & \begin{bmatrix} \text{DAYNAME} & \textit{D4} \\ \text{DAYNUM} & \textit{13} \end{bmatrix} \\ \text{MONTH} & \textit{05} \end{bmatrix} \end{bmatrix} \end{bmatrix} \end{bmatrix}$$

A Compact Encoding of Local Semantics

- Aim:
  - To provide a representation close to the existing ISO format used in TIMEX2
- Benefits:
  - Reuse of existing evaluation tools
  - Easier to read and work with than feature structures
- Two key areas:
  - Underspecification
  - Relative Specification

Underspecification

- In the AVM representation, underspecification is represented by absence of an attribute-value pair
- In our compact encoding, represented by the presence of a lowercase x:
  - <TIMEX TVAL="xxxx-05-13">13th May</TIMEX>
- For ambiguous time we use lowercase t:
  - I expect to see you at <TIMEX T-VAL="xxxx-xx-xxt08:30"> half past eight </TIMEX>.

Underspecified Dates and Times

String	Representation
9 pm	xxxx-xx-xxT21
11:59 pm	xxxx-xx-xxT23:59
eleven in the morning	xxxx-xx-xxT11:00
ten minutes to 3	xxxx-xx-xxt02:50
15 minutes after the hour	xxxx-xx-xxtxx:15
the nineteenth	xxxx-xx-19
January 3	xxxx-01-03
November	xxxx-11
summer	xxxx-SU
’63	xx63
the ’60s	xx6

Relative Specification: Dates

Represented by the prefixes '+' and '–'

String	Representation
today	+0000-00-00
tomorrow	+0000-00-01
yesterday	−0000-00-01
five days ago	−0000-00-05
last month	−0000-01
last summer	−0001-SU
two weeks ago	−0000-W02
this weekend	+0000-W00-WE
this year	+0000
three years ago	−0003
the next century	+01

Relative Specification: Times

String	Representation
sixty seconds later	+0000-00-00T+00:00:60
five minutes ago	+0000-00-00T−00:05
in six hours time	+0000-00-00T+06:00
at 6 a.m. today	+0000-00-00T06:00
last night	−0000-00-01TNI

Cyclic Calendar Dates: Weekdays

- Example 1:
  - We left on Tuesday.
  - T-VAL = "D2"
  - VAL = "2006-04-11"
- Example 2:
  - We left on Tuesday morning.
  - T-VAL = "D2TMO"
  - VAL = "2006-04-11TMO"

Relative References to Days and Months

String	Representation
last Monday	<D1
next Wednesday	>D3
last March	<M03
next March	>M03

Interpretation of Underspecified Values

- Achieved via unification with the reference date.
- Reference date = document creation date or another date from the context:
  - We expect your reply within two days from now.
  - On Monday, we opened the box. A day later, we looked inside. Then, the next day, we took out the contents. Two days after that, we put them back again.
- Takes into account tense used:
  - On Thursday I will give a talk.
  - On Thursday I gave a talk.

Interpretation of Relative Values

- Value calculated on the basis of the reference date:
  - <TIMEX T-VAL="-0000-00-02">Two days ago</TIMEX> we went to Ikea.
  - <TIMEX T-VAL="+0000-00-02">In two days</TIMEX> we will go to Ikea.
- Calendar arithmetic required for some relative expressions:
  - <TIMEX T-VAL=">D1">next Monday</TIMEX>
  - <TIMEX T-VAL="<M03">last March</TIMEX>

Agenda

- Background
- The DANTE System and Its Architecture
- Local Semantics
- Evaluation
- Error Analysis
- Future Work

Results on ACE 2007 Evaluation Corpus

	Unweighted			Weighted			
Domain	P	R	F	P	R	F	ACE Score
Broadcast Conversations	47.9	49.3	48.6	65.0	61.6	63.3	46.5
Broadcast News	55.6	66.8	60.6	70.5	74.7	72.5	55.2
Newswire	56.0	57.2	56.6	69.7	68.5	69.1	58.8
Telephone Conversations	41.5	48.6	44.7	66.1	66.5	66.3	51.4
Usenet	61.8	61.1	61.4	77.6	72.2	74.8	65.3
Weblogs	53.3	54.5	53.9	68.1	68.0	68.0	57.3
Overall	54.7	57.6	56.1	69.7	69.2	69.4	57.2

Results for Normalisation

TIMEX2 Attribute	Precision	Recall	F-Measure
VAL	99.8%	98.0%	98.9%
MOD	76.0%	75.0%	75.5%
SET	100.0%	100.0%	100.0%
ANCHOR_VAL	88.4%	83.5%	85.9%
ANCHOR_DIR	88.1%	87.4%	87.8%

Agenda

- Background
- The DANTE System and Its Architecture
- Local Semantics
- Evaluation
- Error Analysis
- Future Work

Error Analysis for Recognition

- 1056 spurious matches (51.09%)
  - ~50% are missing in the gold standard
  - ~50% are ambiguous: now, fall, a second, may, march
- 586 missing expressions (28.35%)
  - Many based on time trigger word
  - Event based expressions
- 425 extent errors (20.56%)
  - Missing time zone information
  - Modified expressions
  - Expressions built from smaller constituent expressions

Error Analysis for Interpretation

- Gold standard data set contains 5428 TIMEX2 annotations
- Incorrect results:
  - 1460 for the VAL attribute
  - 1067 for the ANCHOR VAL attribute
  - 897 for the ANCHOR DIR attribute
  - 192 for the MOD attribute
  - 53 for the SET attribute
- Note that many of the gold standard annotations are incorrect (eg, approximately 15% of weekdays are wrongly annotated)

Agenda

- Background
- The DANTE System and Its Architecture
- Local Semantics
- Evaluation
- Error Analysis
- Future Work

Future Work

1. Further development of the recognition grammar
2. Improvements to mechanisms in the interpretation module, particularly to the tracking of temporal focus
3. Event time-stamping