Natural Language Processing
SoSe 2015

Relation Extraction

Dr. Mariana Neves

June 15th, 2015

(based on the slides of Dr. Saeedeh Momtazi)

Outline

- Introduction
- Task
- Approaches
  - Pattern Extraction
  - Supervised Learning
  - Semi-supervised Learning
- Temporal and event processing
- Template Filling

Outline

- Introduction
- Task
- Approaches
  - Pattern Extraction
  - Supervised Learning
  - Semi-supervised Learning
- Temporal and event processing
- Template Filling

Information Extraction

The **Hasso Plattner Institute (Hasso-Plattner-Institut für Softwaresystemtechnik GmbH)**, shortly **HPI**, is a German information technology university college, affiliated to the University of Potsdam and is located in Potsdam-Babelsberg nearby Berlin. Teaching and Research of HPI is focused on "IT-Systems Engineering". HPI was founded in 1998 and is the first, and still the only entirely privately funded university college in Germany. It is financed entirely through private funds donated by its founder, Prof. Dr. h.c. Hasso Plattner, who co-founded the largest European software company SAP SE, and is currently the chairman of SAP's supervisory board. President and CEO of HPI is Prof. Dr. Christoph Meinel.[3]

History [edit]

The HPI was founded in 1998 as a public-private partnership. The private partner is the "Hasso Plattner Foundation for Software Systems Engineering", which is the administrative body responsible for the HPI and its only corporate member. The foundation's legal status is that of a GmbH, a limited-liability company according to German law. As the public part of the partnership, the Bundesland Brandenburg provided the estate where several multi-storey buildings were built to form a nice campus. Hasso Plattner declared to provide at least 200 million Euros for the HPI within the first 20 years.[4] He is also actively involved as a lecturer and head of the chair on Enterprise Platforms,[5] where the in-memory technology was developed. In 2004 he received his honorary professorship from the University of Potsdam.

Hasso Plattner Institute

Hasso-Plattner-Institut für Softwaresystemtechnik GmbH

IT Systems Engineering | Universität Potsdam

Motto	Design IT. Create Knowledge.
Established	1998[1]
Type	Private university institute
Director	Prof. Dr. Christoph Meinel
Administrative staff	60[2]
Students	about 480[2]
Location	Potsdam, Germany
Campus	Griebnitzsee
Colors	Orange Vivid orange Dark pink
Affiliations	University of Potsdam
Website	www.hpi.uni-potsdam.de

(http://en.wikipedia.org/wiki/Hasso_Plattner_Institute)

Named Entity Recognition

- HPI is affiliated to the Potsdam University and located in Potsdam near Berlin. It was founded in 1998 by Hasso Plattner, one of the co-founders of the European software company, SAP AG.
  - HPI (ORG)
  - Potsdam University (ORG)
  - Potsdam (LOC)
  - Berlin (LOC)
  - 1998 (DATE)
  - Hasso Plattner (PER)
  - SAP AG (ORG)

Relation Extraction

- HPI is affiliated to the Potsdam University and located in Potsdam near Berlin. It was founded in 1998 by Hasso Plattner, one of the co-founders of the European software company, SAP AG.
  - HPI – Potsdam: located (ORG-LOC)
  - HPI – Berlin: near (ORG-LOC)
  - Potsdam – Berlin: near (LOC-LOC)
  - HPI – 1998: founded (ORG-DATE)
  - HPI - Hasso Plattner: founder (ORG-PER)
  - SAP AG - Hasso Plattner: co-founder (ORG-PER)

Motivation

- Creating new structured data sources (knowledge bases)
  - DBPedia
  - Freebase
  - Yago
  - Infobox in Wikipedia

Motivation

- Answering complex questions using multiple sources
  - Which soccer player married a Spice Girls star?

    ("?x" is-a "soccer player")

    ("?x" married "?y")

    ("?y" member "Spice Girls")

Outline

- Introduction
- Task
- Approaches
  - Pattern Extraction
  - Supervised Learning
  - Semi-supervised Learning
- Temporal and event processing
- Template Filling

Relation Representation

- Representing data as triples
  - (Argument1 RelationType Argument2)
  - (Subject Predicate Object)

("Messi" is-a "soccer player")

("Brad Pitt" married "Angelina Jolie")

("Messi" member "Barcelona FC")

Relation Representation

- Resource Description Framework (RDF)

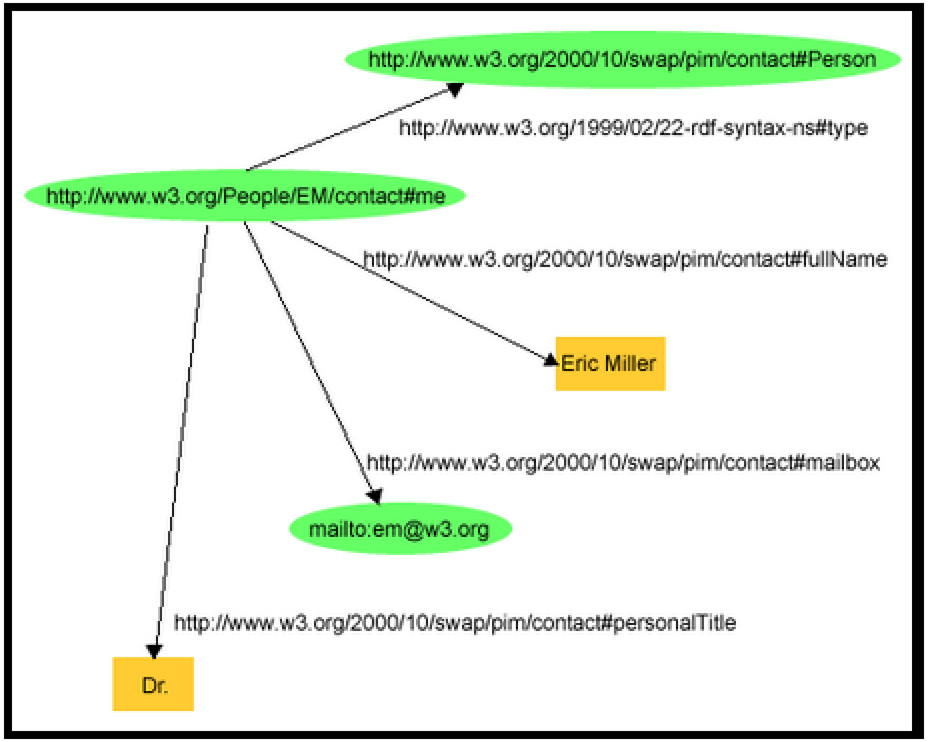

(http://en.wikipedia.org/wiki/Resource_Description_Framework#/media/File:Rdf_graph_for_Eric_Miller.png)

Relation Types

- Having various relation types based on the type of arguments
  - PER-PER: Spouse, Parent, Child, Friendship, Colleague, …

("Brad Pitt" married "Angelina Jolie")

("Shiloh Nouvel Jolie-Pitt" child "Angelina Jolie")

("Messi" colleague "Neymar")

Relation Types

- Having various relation types based on the type of arguments
  - PER-LOC: Place of birth, Lives in, Place of death, Buried in, ...

("Angela Merkel" place_of_birth "Hamburg")

("Angela Merkel" lives "Berlin")

("Beethoven" place_of_birth "Bonn")

("Beethoven" place_of_death "Vienna")

("Beethoven" buried "Vienna")

Relation Types

- Having various relation types based on the type of arguments
  - PER-ORG: Founder, Co-founder, Owner, Employee, Student/Alum, Professor, …

("Prof. Plattner" founder "HPI")

("Prof. Naumann" professor "HPI")

("Dr. Neves" employee "HPI")

Relation Types

- Having various relation types based on the type of arguments
  - ORG-LOC: Located, Near, Founded-location, Headquarter, ..

("HPI" located "Potsdam")

("Potsdam" near "Berlin")

("Potsdam" headquarter "HPI")

Relation Types

- Having various relation types based on the type of arguments
  - PER-DATE: Date of Birth, Date of Marriage, Date of Death, …

("HPI" founded "1998")

("Angela Merkel" date_of_birth "17-Jul-1954")

("Beethoven" date_of_death "17-Dec-1770")

Approaches

- Manually created patterns
- Supervised machine learning
- Semi-supervised learning

Outline

- Introduction
- Task
- Approaches
  - Pattern Extraction
  - Supervised Learning
  - Semi-supervised Learning
- Temporal and event processing
- Template Filling

Pattern Extraction

- What are the potential words to express a relation type?
  - (PER Member ORG)
  - ("?x" Member "?y")

  - x is a member of y.
  - x is an employee of y.
  - x works at y.
  - x is a staff of y.
  - ...

  - x is (a|an) (member|employee|staff|professor|researcher|lecturer) of y.
  - x (works) at y.

Pattern Extraction

- Advantages
  - Having high precision results
- Disadvantages
  - Having low recall
  - Finding all possible patterns is labor intensive
  - Covering all relations is very difficult
  - Language is complex

Outline

- Introduction
- Task
- Approaches
  - Pattern Extraction
  - Supervised Learning
  - Semi-supervised Learning
- Temporal and event processing
- Template Filling

Supervised Classification

- Training data
  - Define a set of relation types
  - Choosing the corresponding named entities
  - Selecting a set of texts as training data
  - Recognizing the named entities in the text
  - Labeling the relations between named entities manually

Task

- Input
  - A pair of entities (NER)
  - A context in which this pair appears
  - Possible relation types
- Output
  - Type of relation between two entities, if there exist any

Task

- „Thomas Edison died on October 18, 1931, in New Jersey due to complications of diabetes."
  - PER-LOC (Thomas Edison, New Jersey)
  - Place of birth, Place of death, Buried in

Feature Selection

- „Thomas Edison died on October 18, 1931, in New Jersey due to complications of diabetes."
- The target entities
  - T1: Thomas Edison
  - T2: New Jersey

Feature Selection

- „Thomas Edison died on October 18, 1931, in New Jersey due to complications of diabetes.“
- The named entity label of the target words (blind entities)
  - NE(T1): PER
  - NE(T2): LOC

Feature Selection

- „Thomas Edison died on October 18, 1931, in New Jersey due to complications of diabetes.“
- Bag-of-words
  - 1931 October died 18 , on , in
- Bag-of-bigrams
  - [1931 ,] [October 18] [died on] [18 ,] [, 1931] [on October] [, in]

Feature Selection

- „Thomas Edison died on October 18, 1931, in New Jersey due to complications of diabetes."
- Bag-of-words, entities
  - YEAR MONTH died DATE , on , in
- Bag-of-bigrams, entities
  - [YEAR ,] [MONTH DATE] [died on] [DATE ,] [, YEAR] [on MONTH] [, in]

Feature Selection

- „Thomas Edison died on October 18, 1931, in New Jersey due to complications of diabetes."
- Bag-of-words, entities, stems
  - YEAR MONTH die DATE , on , in
- Bag-of-bigrams, entities, stems
  - [YEAR ,] [MONTH DATE] [die on] [DATE ,] [, YEAR] [on MONTH] [, in]

Feature Selection

- „Thomas Edison died on October 18, 1931, in New Jersey due to complications of diabetes."
- Distance in words between arguments
  - 6 words
  - 8 words (w/ punctuations)

Feature Selection

- „Thomas Edison died on October 18, 1931, in New Jersey due to complications of diabetes."
- Number of entities between arguments
  - None?
  - Three (MONTH, DATE, YEAR)

Feature Selection

- „Thomas Edison died on October 18, 1931, in New Jersey due to complications of diabetes."
- Surrounding words of target entities
  - For instance, [-1,+1]
    - $T1_{+1}$: died
    - $T2_{-1}$: in
    - $T2_{+1}$: due

Feature Selection

- „Thomas Edison died on October 18, 1931, in New Jersey due to complications of diabetes."
- Bags of chunk heads
  - VP PP NP

Word	Base Form	Part-Of-Speech	Chunk
Thomas/NP	Thomas/NP	NNP	B-NP
Edison/NP	Edison/NP	NNP	I-NP
died/VP	died/VP	NN	I-NP
on	on	IN	B-PP
October	October	NNP	B-NP
18	18	CD	I-NP
,	,	,	I-NP
1931	1931	CD	I-NP
,	,	,	O
in	in	IN	B-PP
New	New	NNP	B-NP
Jersey	Jersey	NNP	I-NP
due	due	IN	B-PP
to	to	TO	B-PP
complications	complication	NNS	B-NP
of	of	IN	B-PP
diabetes	diabete	NNS	B-NP
.	.	.	O
#	#	#	B-NP
#	#	#	I-NP
#	#	#	I-NP
This	This	DT	B-NP
is	be	VBZ	B-VP
a	a	DT	B-NP
sample	sample	NN	I-NP
.	.	.	O
Replace	Replace	VB	B-VP
this	this	DT	B-NP
with	with	IN	B-PP
your	your	PRP$	B-NP
own	own	JJ	I-NP
text	text	NN	I-NP
.	.	.	O

(http://nactem7.mib.man.ac.uk/geniatagger/)

Feature Selection

- „Thomas Edison died on October 18, 1931, in New Jersey due to complications of diabetes.“
- Chunk base-phrase paths
  - VP→PP→NP→NP→NP→O→PP

Word	Base Form	Part-Of-Speech	Chunk
Thomas/NP	Thomas/NP	NNP	B-NP
Edison/NP	Edison/NP	NNP	I-NP
died/VP	died/VP	NN	I-NP
on	on	IN	B-PP
October	October	NNP	B-NP
18	18	CD	I-NP
,	,	,	I-NP
1931	1931	CD	I-NP
,	,	,	O
in	in	IN	B-PP
New	New	NNP	B-NP
Jersey	Jersey	NNP	I-NP
due	due	IN	B-PP
to	to	TO	B-PP
complications	complication	NNS	B-NP
of	of	IN	B-PP
diabetes	diabete	NNS	B-NP
.	.	.	O
#	#	#	B-NP
#	#	#	I-NP
#	#	#	I-NP
This	This	DT	B-NP
is	be	VBZ	B-VP
a	a	DT	B-NP
sample	sample	NN	I-NP
.	.	.	O
Replace	Replace	VB	B-VP
this	this	DT	B-NP
with	with	IN	B-PP
your	your	PRP$	B-NP
own	own	JJ	I-NP
text	text	NN	I-NP
.	.	.	O

(http://nactem7.mib.man.ac.uk/geniatagger/)

Feature Selection

- „Thomas Edison died on October 18, 1931, in New Jersey due to complications of diabetes."
- Constituent-tree paths
  - VP→PP→NP

```
(ROOT
  (S
    (NP (NNP Thomas) (NNP Edison))
    (VP (VBD died)
      (PP (IN on)
        (NP (NNP October) (CD 18) (, ,) (CD 1931) (, ,)))
      (PP (IN in)
        (NP
          (NP (NNP New) (NNP Jersey))
          (ADJP (JJ due)
            (PP (TO to)
              (NP
                (NP (NNS complications))
                (PP (IN of)
                  (NP (NN diabetes)))))))))
    (. .)))
```

(http://nlp.stanford.edu:8080/parser/index.jsp)

Feature Selection

- „Thomas Edison died on October 18, 1931, in New Jersey due to complications of diabetes.“
- Dependency-tree paths
    - nsubj-nmod
    - compound-nsubj-nmod-compound

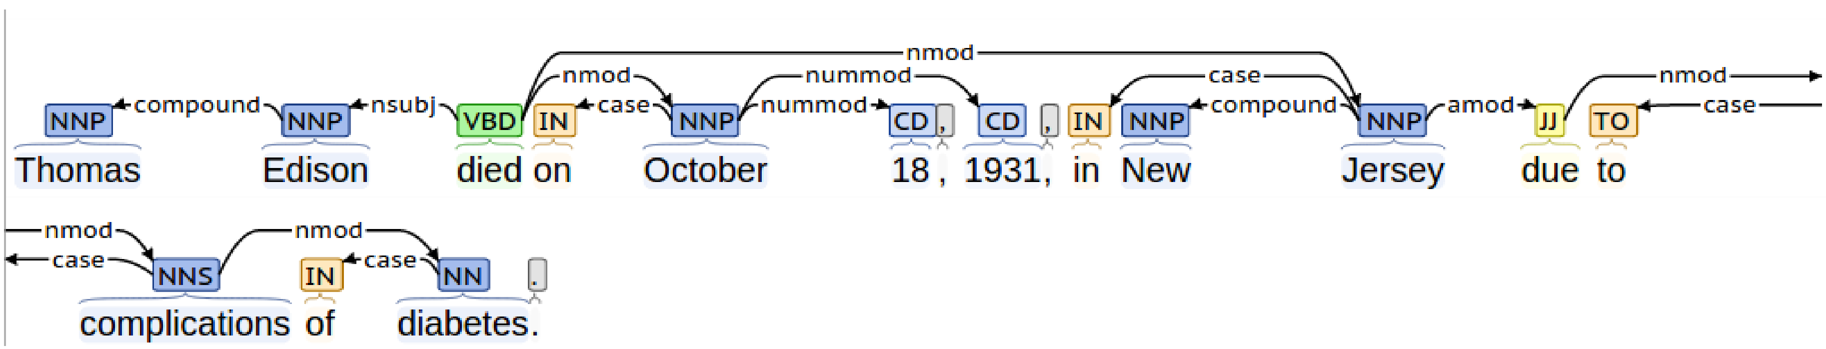

(http://nlp.stanford.edu:8080/corenlp/process)

Feature Selection

- „Thomas Edison died on October 18, 1931, in New Jersey due to complications of diabetes.“
- Tree distance between arguments
    - Two (nsubj-nmod)
    - Four (compound-nsubj-nmod-compound)

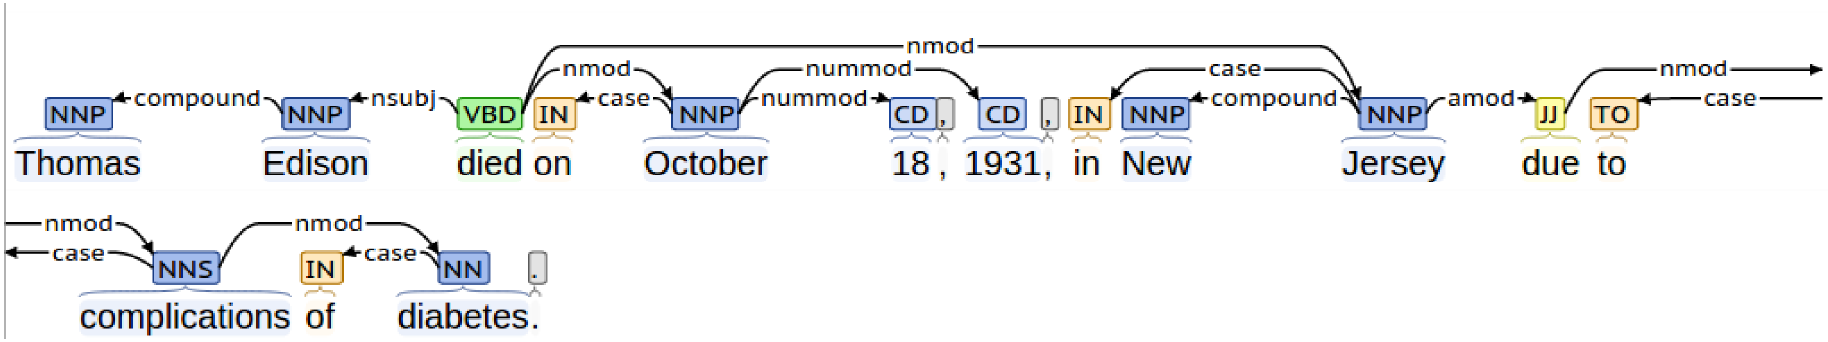

(http://nlp.stanford.edu:8080/corenlp/process)

Classification Algorithm

- Applying any of the classifiers
  - K Nearest Neighbor
  - Support Vector Machines
  - Naïve Bayes
  - Maximum Entropy
  - Logistic Regression
  - ...

Supervised Classification

- Advantages
  - Very good performance if
    - enough training data
    - test data similar to training data
- Disadvantages
  - Manual labeling of training data is labor expensive
  - Difficult to get good results for other domains and relations

Outline

- Introduction
- Task
- Approaches
  - Pattern Extraction
  - Supervised Learning
  - Semi-supervised Learning
- Temporal and event processing
- Template Filling

Semi-supervised Learning

- Having no large training data
  - but a large collection of documents
- Producing a small training data (seed data)
  - A set of triples
- Bootstrapping
  - Using the seed data to find further entity pairs with the same relation

Bootstrapping

- Using the collected seed data
- Finding sentences which contain at least one entity pair
- Extracting the common contexts of the pair
- Creating patterns (or models) from the extracted context
- Using the pattern (or model) to get more pairs and add them to seed data

Bootstrapping

- Using the collected seed data
  - (Thomas Edison Spouse Mina Mille)
  - (Brad Pitt Spouse Angelina Jolie)
  - ...

Bootstrapping

- Using the collected seed data
- Finding sentences which contain at least one entity pair

- Thomas Edison married Mina Mille.
- Edison married a young woman named Mina Mille.
- In 1871, Thomas Edison married Mina Mille.
- Thomas Edison marries Mina Mille on December 25.

Bootstrapping

- Using the collected seed data
- Finding sentences which contain at least one entity pair
- Extracting the common contexts of the pair
- Creating patterns (or models) from the extracted context

- Thomas Edison married Mina Mille.
- Edison married a young woman named Mina Mille.
- In 1871, Thomas Edison married Mina Mille.
- Thomas Edison marries Mina Mille on December 25.

Bootstrapping

- Using the collected seed data
- Finding sentences which contain at least one entity pair
- Extracting the common contexts of the pair
- Creating patterns (or models) from the extracted context
- Using the pattern (or model) to get more pairs and add them to seed data
  - (Albert Einstein Spouse "?")

Bootstrapping

- Einstein marries his cousin Elsa Löwenthal on June 2.
- Einstein married Elsa Löwenthal in Berlin.
- Einstein married Elsa Löwenthal on 2 June 1919.
- After their divorce in 1919, Einstein married Elsa Löwenthal in the same year.
- Albert Einstein was married to Elsa Löwenthal for 17 years.
- Einstein marries Elsa Löwenthal.
- In the same year Albert Einstein married Elsa Löwenthal.

⇒ (Albert Eistein Spouse Elsa Löwenthal)

Bootstrapping

- Using the collected seed data (start over again)
  - (Thomas Edison Spouse Mina Mille)
  - (Brad Pitt Spouse Angelina Jolie)
  - ...
  - (Albert Eistein Spouse Elsa Löwenthal)

Bootstrapping

- Using the collected seed data
- Finding sentences which contain at least one entity pairs
- Extracting the common contexts of the pair
- Creating patterns (or models) from the extracted context

- Albert Einstein's wife, Elsa Löwenthal, was his first cousin.
- Elsa Löwenthal was the wife of Albert Einstein.
- Einstein's wife was named Elsa Löwenthal.

Semantic drift

- Erroneous patterns → introduction of erroneous tuples → problematics patterns

- Brad Pitt married the daughter of Jon Voigth

Assessment of patterns

- Assess new pattern (p)
  - regarding current set of tuples (T)
  - regarding produtivity in the document collection (D)

Assessment of patterns

- (Riloff and Jones 1999)

$$Conf_{RlogF}(p) = \frac{hits_p}{hits_p + misses_p} \cdot \log(finds_p)$$

- *hits*: set of tuples in T that p <u>matches</u> while looking in D
- *misses*: set of tuples in T that p <u>misses</u> while looking at D
- *finds*: total set of tuples that p finds in D

Outline

- Introduction
- Task
- Approaches
  - Pattern Extraction
  - Supervised Learning
  - Semi-supervised Learning
- Temporal and event processing
- Template Filling

Temporal and event processing

- Analyze events and their relations to each other
- Applications
  - Summarization
  - Question answering

When did airlines as a group last raise fares?

Last week, Delta boosted thousands of fares by $10 per round trip, and most big rivals immediately matched the increase.

(Dateline 7/2/2007)

Temporal and event processing

- Extract temporal expression (last week, immediately)

Last week, Delta boosted thousands of fares by $10 per round trip, and most big rivals immediately matched the increase.

(Dateline 7/2/2007)

Temporal and event processing

- Figure out the time the expressions refer

Last week, Delta boosted thousands of fares by $10 per round trip, and most big rivals immediately matched the increase.

(Dateline 7/2/2007)

[the week before 7/2/2007]

Temporal and event processing

- Detecting events

Last week, Delta <u>boosted thousands of fares by $10</u> per round trip, and most big rivals immediately matched the increase.

(Dateline 7/2/2007)

[Fares raise]

[Rivals matched increase]

Temporal and event processing

- Associate times with events

Last week, Delta boosted thousands of fares by $10 per round trip, and most big rivals immediately matched the increase.

(Dateline 7/2/2007)

[Fares raise in the week before 7/2/2007]

[Rivals matched increase in the week before 7/2/2007]

Temporal expression recognition

- Absolute points in time (7/2/2007)
- Relative times (last week, immediately)
- Durations (during a week)
- ...

Temporal expression recognition

- Time expression construction is kind of conventionalized

- Rule-based systems
- Statistical sequence classifiers
- Constituent-based classification

Rule-based systems

- Patterns based on syntactic chunks
- Should include temporal lexical triggers

```
#i.e. "3-year","four-month old"
   $matchPatterns[0] = "($OT\+$numString(-|$CT\+\\s$OT\+)$TE_Units((-|(s)?$CT\+\\s$OT\+)old)?$CT\+)";

#i.e. "the past twenty four years"
   $matchPatterns[1] = "($OT\+the$CT\+\\s$OT\+(\[pl\]ast|next)$CT\+\\s$OT\+$numString$CT\+\\s$OT\+$TE_Units(s)?$CT\+)";
```

```
# hh:mm AM/PM (time zone)
if($string =~ /[ap]\.?m\b/io) {
    $string =~ s/(($OT+(about|around|some|exactly|precisely)$CT+\s+)?($OT*(quarter|half|$OTCD$OT*\w+$CT+\s+$OT+minutes?)$CT+\s
    $string =~ s/(($OT+(about|around|some|exactly|precisely)$CT+\s+)?($OT*(quarter|half|$OTCD$OT*((\w+(-\w+)?|\d\d?)$CT+\s+$OT
}
```

(MITRE's TempEx code
https://github.com/gangeli/ParsingTime/blob/master/etc/TempEx.pm)

Statistical sequence classifiers

- Similar to named-entity recognition
- Labeling using IOB tags

A fare increase initiated last/B week/I by Delta Airlines ...

Statistical sequence classifiers

- Features for a machine learning approach
  - Token
  - Bag of words around
  - Shape
  - Part-of-speech tags
  - Chunk tags
  - Lexical triggers (temporal terms)

Constituent-based classification

- Start with automatic parsing

```
(ROOT
  (S
    (S
      (NP (JJ Last) (NN week))
      (, ,)
      (NP (NNP Delta))
      (VP (VBD boosted)
        (NP
          (NP (NNS thousands))
          (PP (IN of)
            (NP (NNS fares))))
        (PP (IN by)
          (NP
            (NP ($ $) (CD 10))
            (PP (IN per)
              (NP (NN round) (NN trip)))))))
    (, ,)
    (CC and)
    (S
      (NP (RBS most) (JJ big) (NNS rivals))
      (ADVP (RB immediately))
      (VP (VBD matched)
        (NP (DT the) (NN increase))))
    (. .)))
```

Constituent-based classification

- Start with automatic parsing
- Classifying each node w.r.t. the presence of temporal terms

```
(ROOT
  (S
    (S
      (NP (JJ Last) (NN week))
      (, ,)
      (NP (NNP Delta))
      (VP (VBD boosted)
        (NP
          (NP (NNS thousands))
          (PP (IN of)
            (NP (NNS fares))))
        (PP (IN by)
          (NP
            (NP ($ $) (CD 10))
            (PP (IN per)
              (NP (NN round) (NN trip)))))))
    (, ,)
    (CC and)
    (S
      (NP (RBS most) (JJ big) (NNS rivals))
      (ADVP (RB immediately))
      (VP (VBD matched)
        (NP (DT the) (NN increase))))
    (. .)))
```

Temporal normalization

- w.r.t. the ISO 8601 standard for encoding temporal values

- Document date: 2007-07-02
- Event date: 2007-W26
- Durations: P1WE, P3Y, P20D

Temporal normalization

- Approaches
  - Rule-based methods
    - Fully qualified temporal expression
    - Absolute temporal expression
    - Relative temporal expression
    - Durations

Event detection and analysis

- Last week, Delta boosted thousands of fares by $10 per round trip, and most big rivals immediately matched the increase.

Event detection

- Approaches
  - Rule-based methods
  - Machine learning

Event analysis

- Order events
  - Partial ordering (binary relation detection)
  - Before, after, during relations

Temporal relations

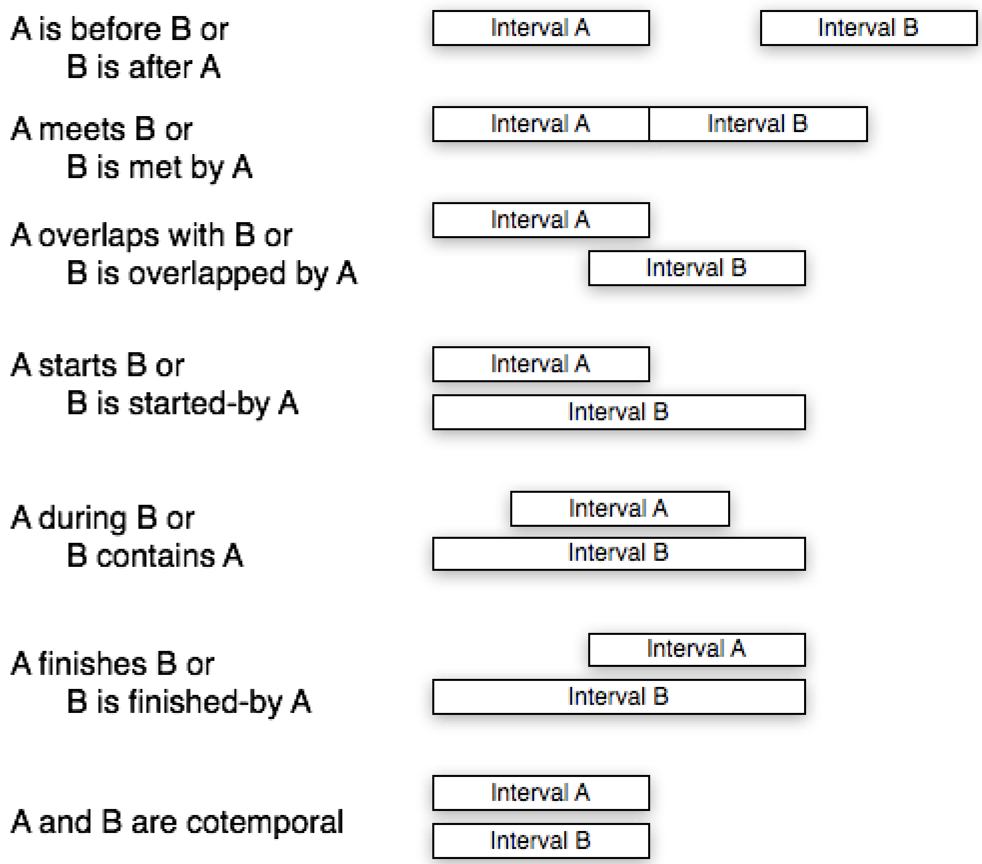

(http://franz.com/agraph/support/documentation/current/allen-relations.png)

Outline

- Introduction
- Task
- Approaches
  - Pattern Extraction
  - Supervised Learning
  - Semi-supervised Learning
- Temporal and event processing
- Template Filling

Template filling

- Template
  - slots

Hasso Plattner Institute

Hasso-Plattner-Institut für Softwaresystemtechnik GmbH

IT Systems Engineering | Universität Potsdam

Motto	Design IT. Create Knowledge.
Established	1998[1]
Type	Private university institute
Director	Prof. Dr. Christoph Meinel
Administrative staff	60[2]
Students	about 480[2]
Location	Potsdam, Germany
Campus	Griebnitzsee
Colors	Orange Vivid orange Dark pink
Affiliations	University of Potsdam
Website	www.hpi.uni-potsdam.de

Statistical template filling

- Train separate classifiers, one for each slot

Name	INSTITUTION	Hasso-Plater Institute
Year foundation	YEAR	1998
Director	PEOPLE	Prof. Meinel
Location	CITY COUNTRY	Potsdam Germany
Affiliation	INSTITUTION UNIVERSITY	University of Potsdam

Statistical template filling

- Train separate classifiers, one for each slot
- Challenges
  - Multiple text segments labeled with the same slot label
    - Christoph Meinel, Prof. Meinel

Statistical template filling

- Train separate classifiers, one for each slot
- Challenges
  - Multiple entities of the expected type for a given slot
    - Potsdam, Germany, Berlin, Haifa, etc.
    - University of Potsdam, Stanford University, Cape Town University, Nanjing University, etc.

Statistical template filling

- Train one large classifier, usually Hidden Markov Model
  - Sequential labeling
  - Potsdam, Berlin, Germany (location) → University of Potsdam (university)

Further Reading

- Speech and Language Processing
  - Chapter 22