Textual inference: Methods, open source platform and applications

Günter Neumann, January, 2014

The following slides are based on a presentation by Dagan, Magnini, Neumann and Pado held at the "Symposium on Semantic Text Processing - Industrial Outlook Bar-Ilan University, Department of Computer Science, NLP Lab, Nov. 18-19 2014", cf. http://u.cs.biu.ac.il/~nlp/workshop14/

They are accompanied with further additional material.

UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

Textual inference:
Methods, open source platform and applications

Ido Dagan

Bar-Ilan University, Israel

Bernardo Magnini

Foundation Bruno Kessler, Trento

Guenter Neumann

German Research Center for Artificial Intelligence, Saabrucken

Sebastian Pado

University of Heidelberg

Excitement project

What is applied textual inference?

What is applied textual inference?

"Match" different text fragments where:

What is applied textual inference?

"Match" different text fragments where:

One text has the **same meaning** as the other

What is applied textual inference?

"Match" different text fragments where:

One text has the **same meaning** as the other

pepper may trigger sneezing

pepper can cause sneezing

What is applied textual inference?

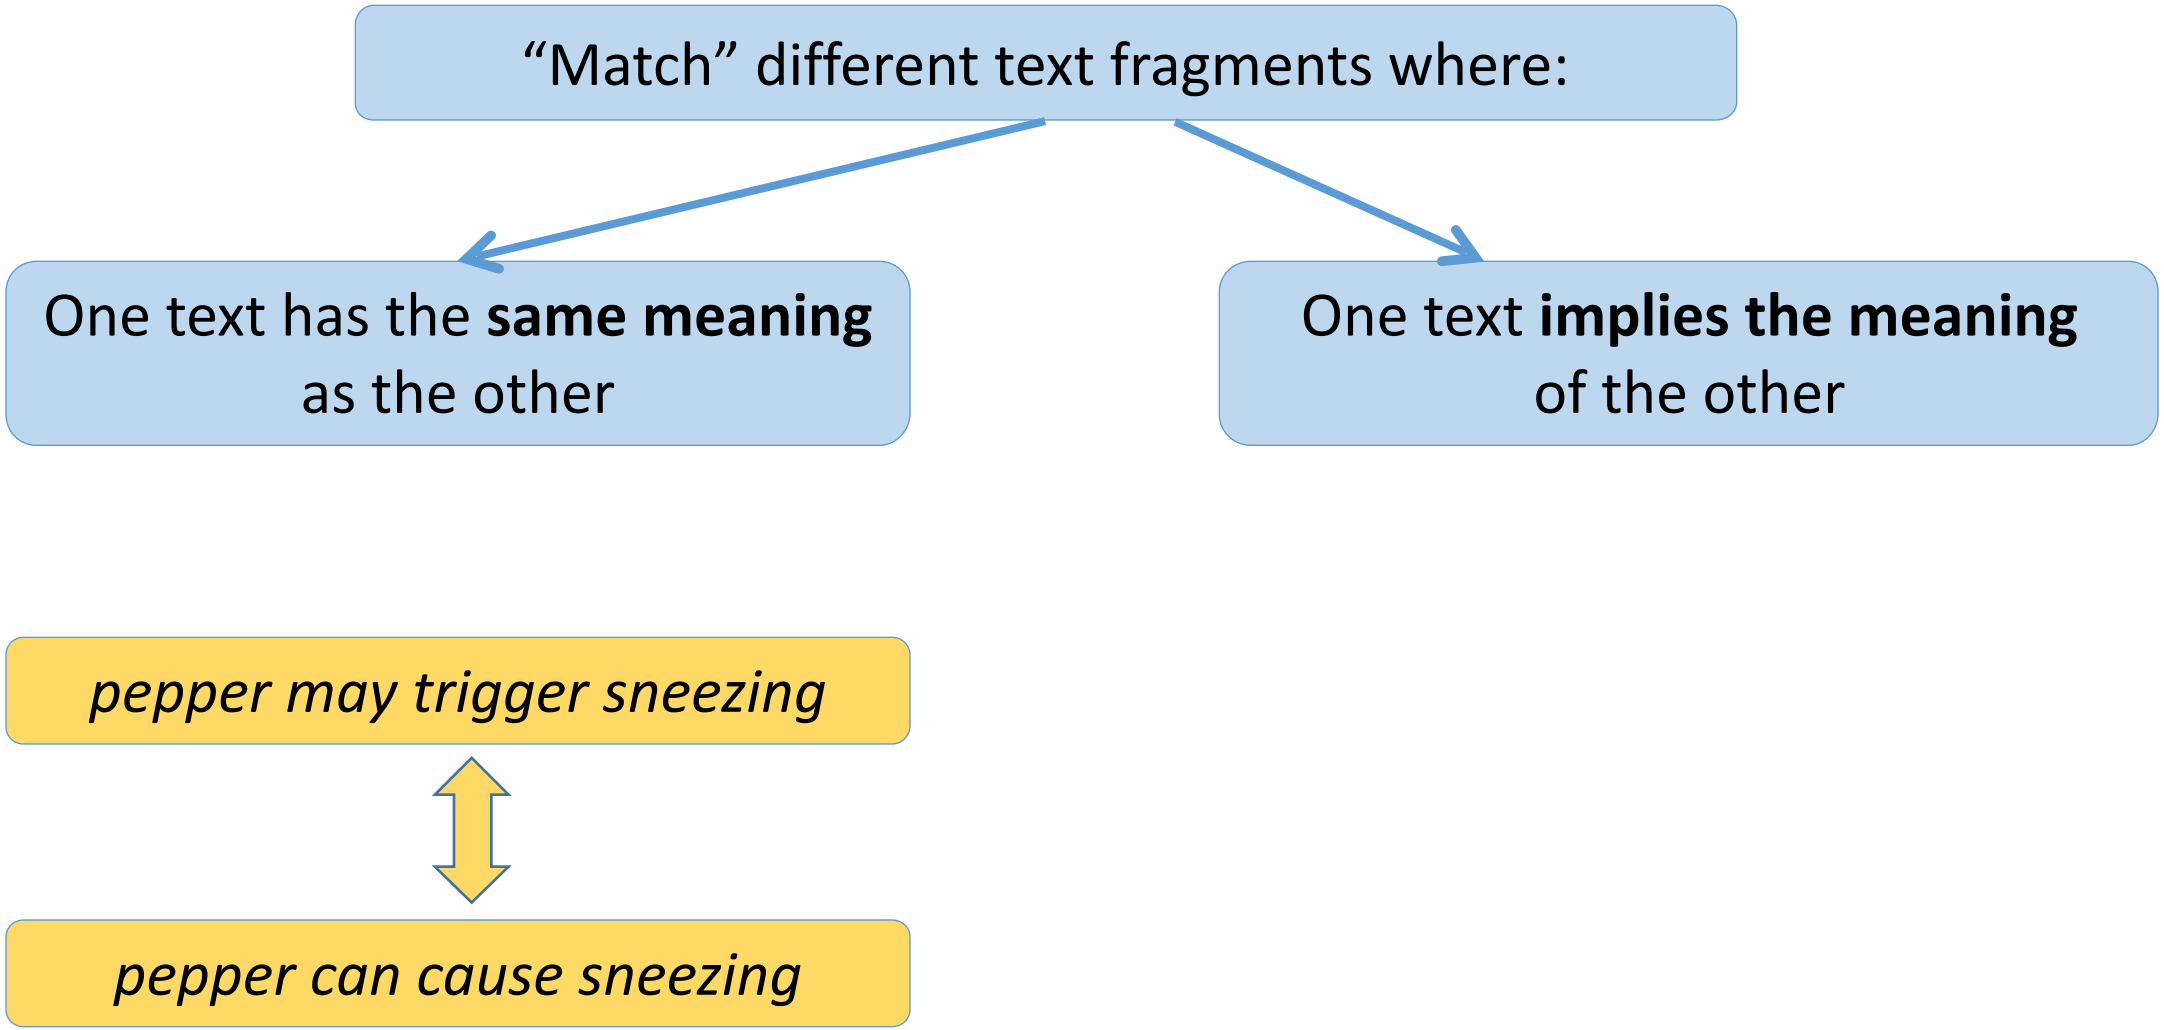

What is applied textual inference?

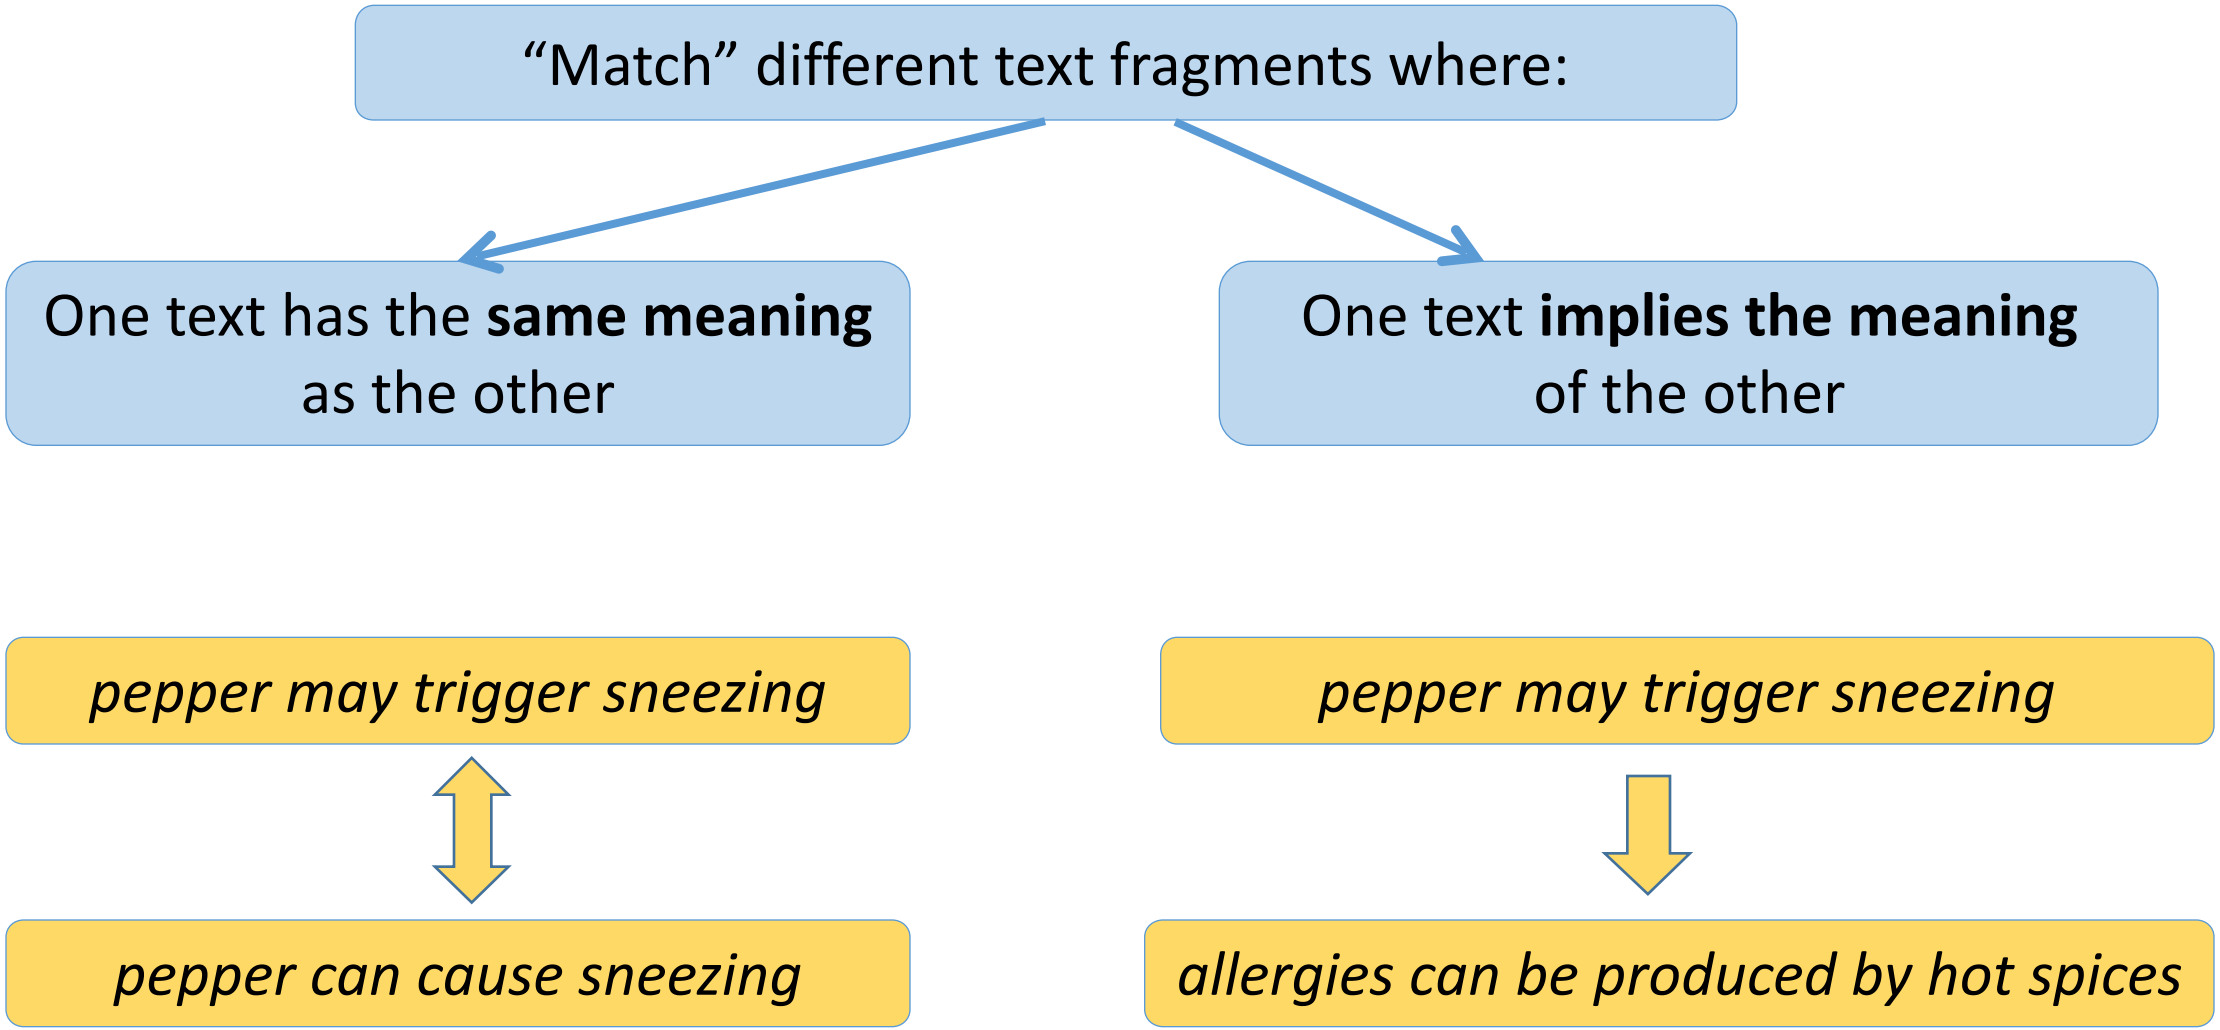

What is applied textual inference?

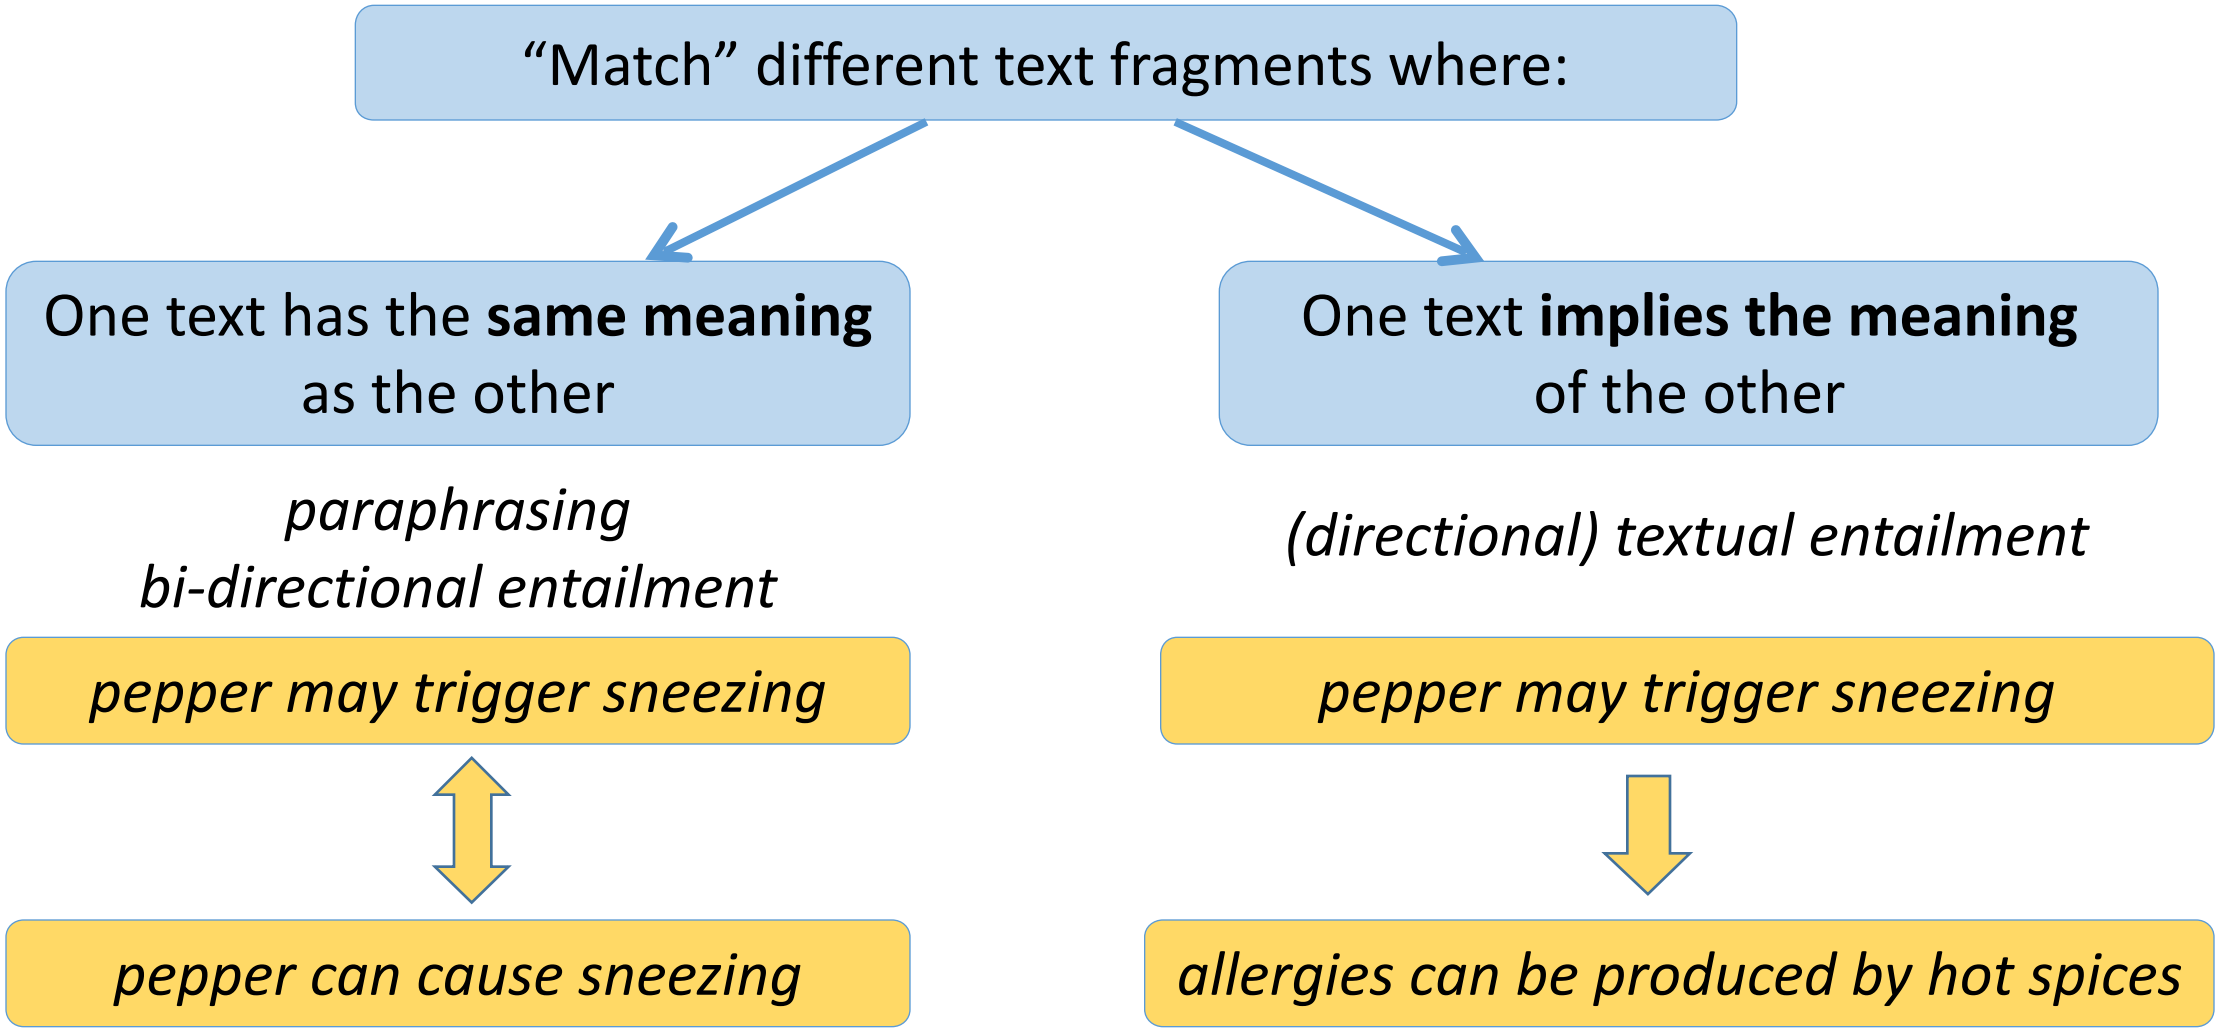

Example Applications

Question Answering

Which foods are allergenic?

Example Applications

Question Answering

Which foods are allergenic?

allergies can be produced by hot spices

Example Applications

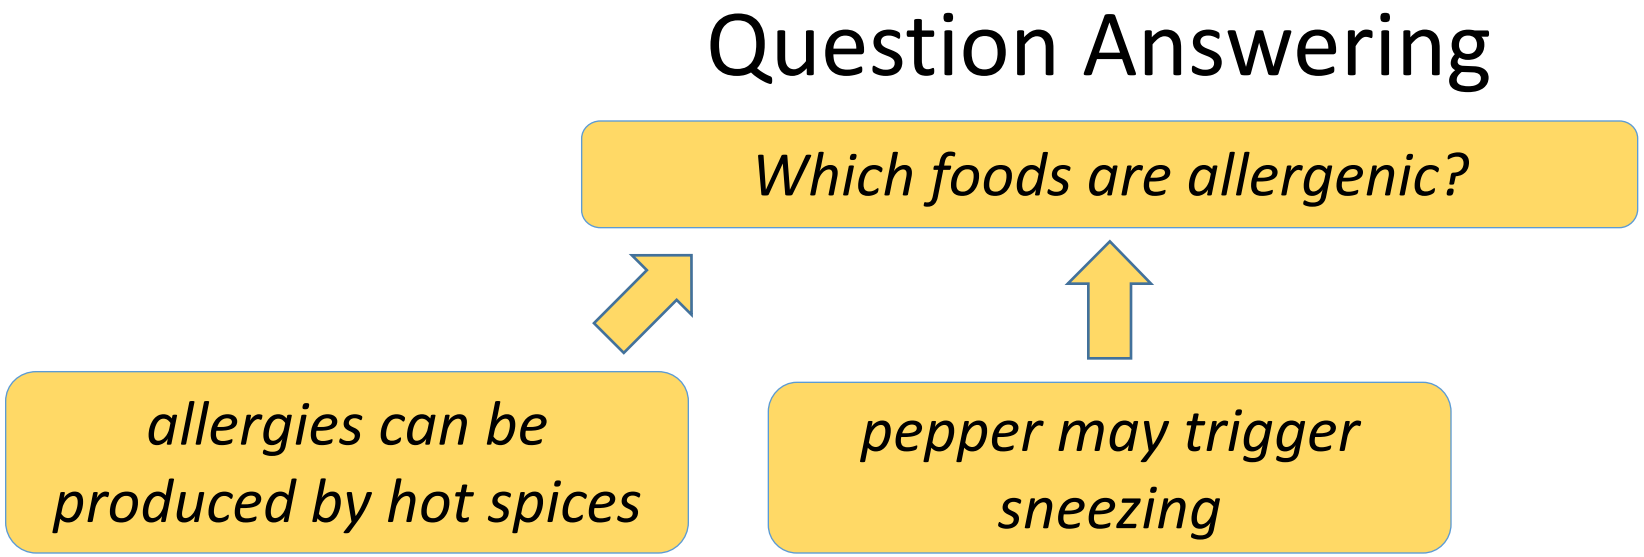

Example Applications

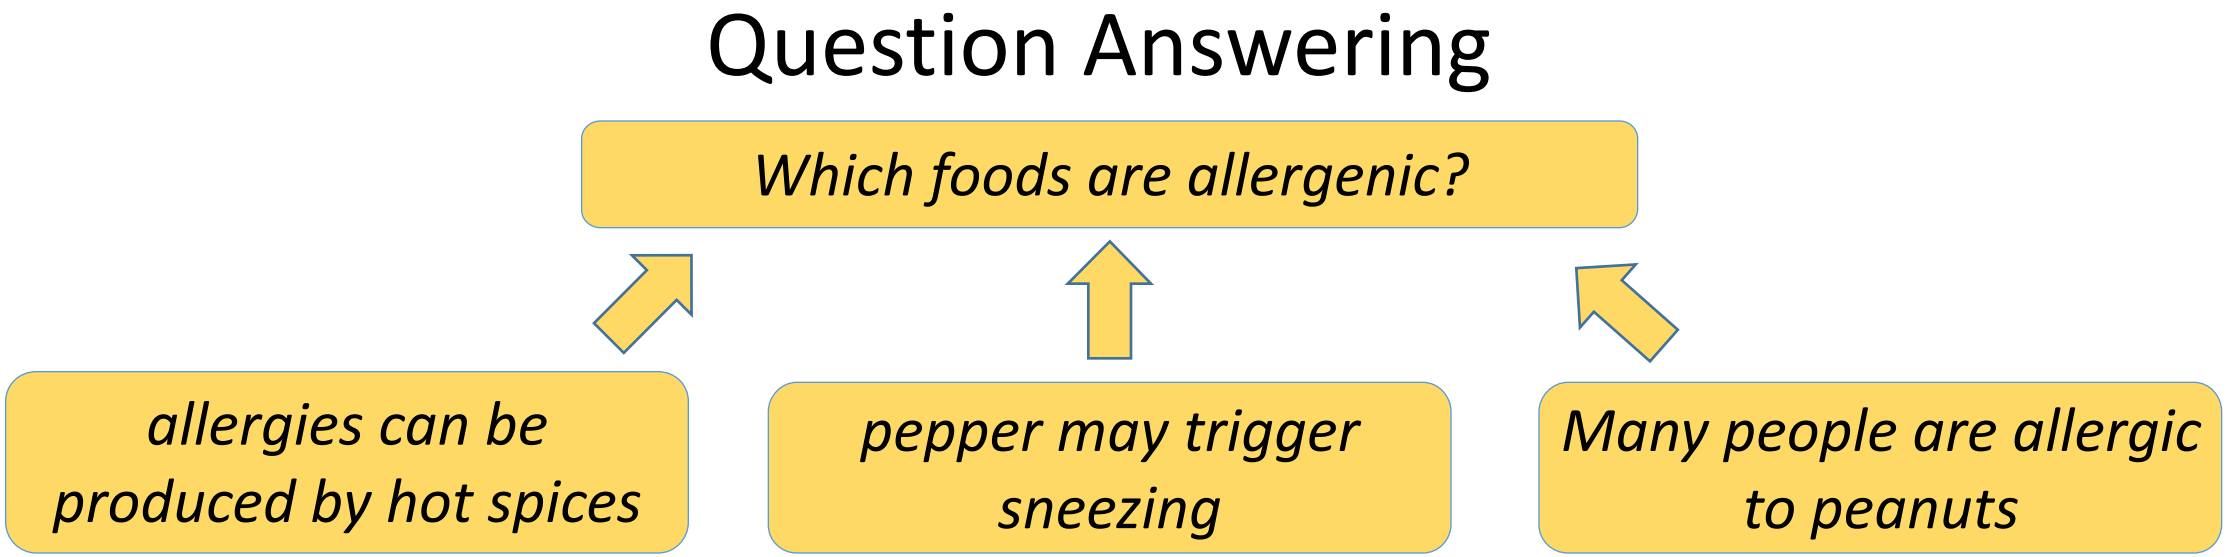

Example Applications

Question Answering

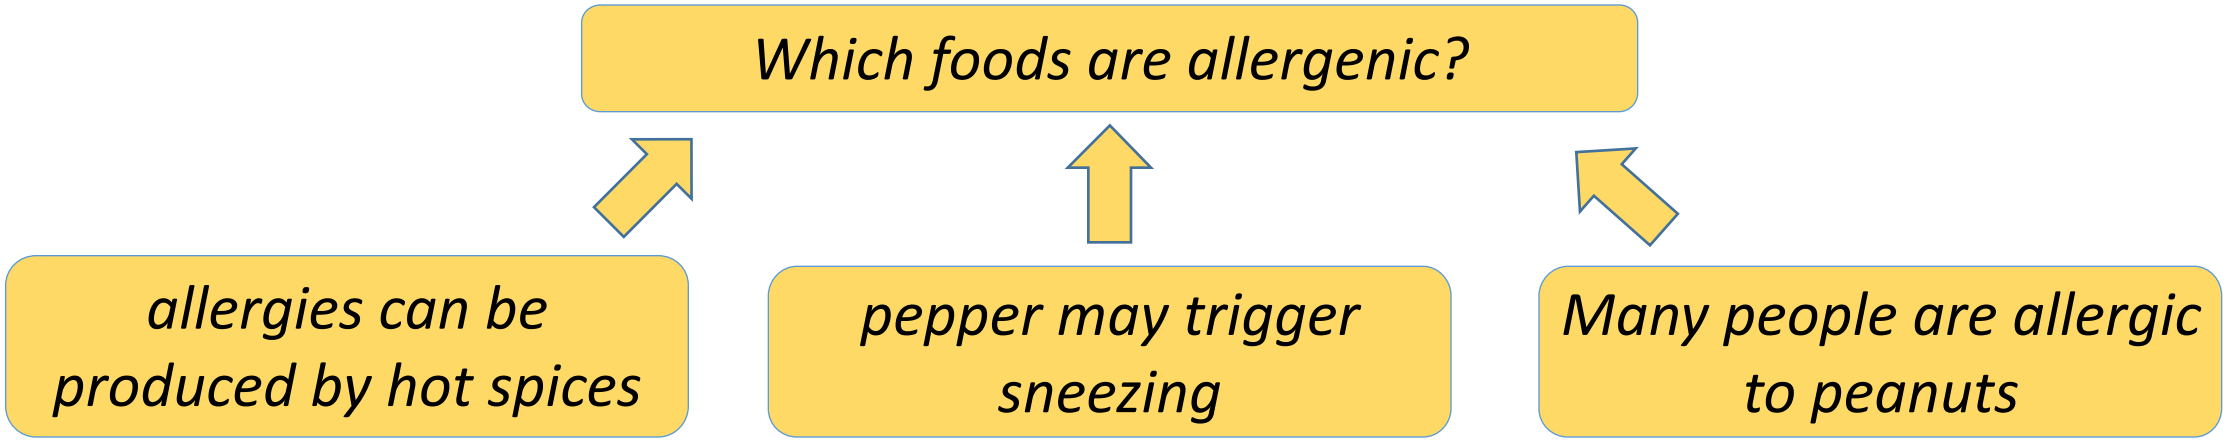

Search

allergenic foods

Example Applications

Question Answering

Which foods are allergenic?

allergies can be produced by hot spices

pepper may trigger sneezing

Many people are allergic to peanuts

Information Extraction

Extract pairs of foods and symptoms

Search

allergenic foods

Example Applications

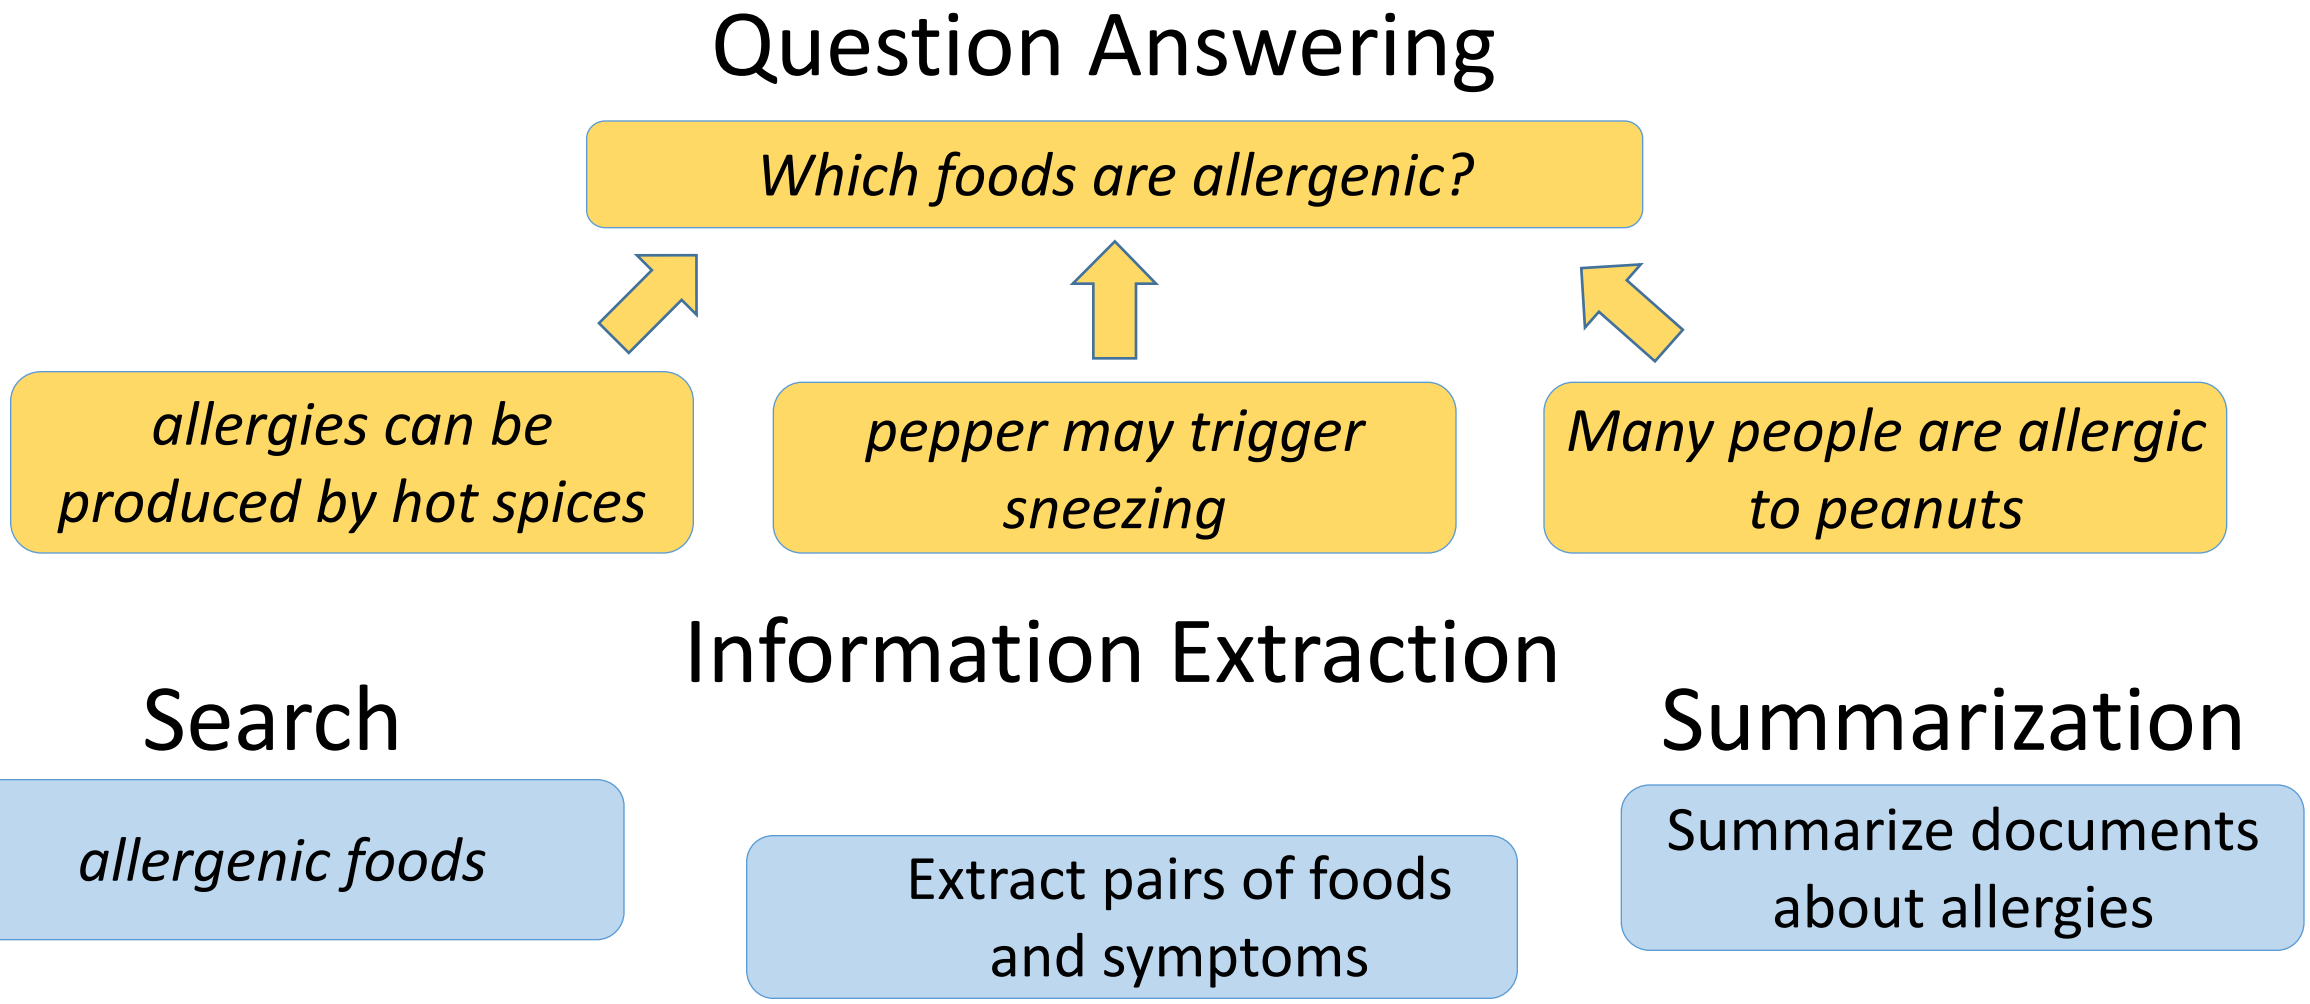

Novel Application: Text Exploration

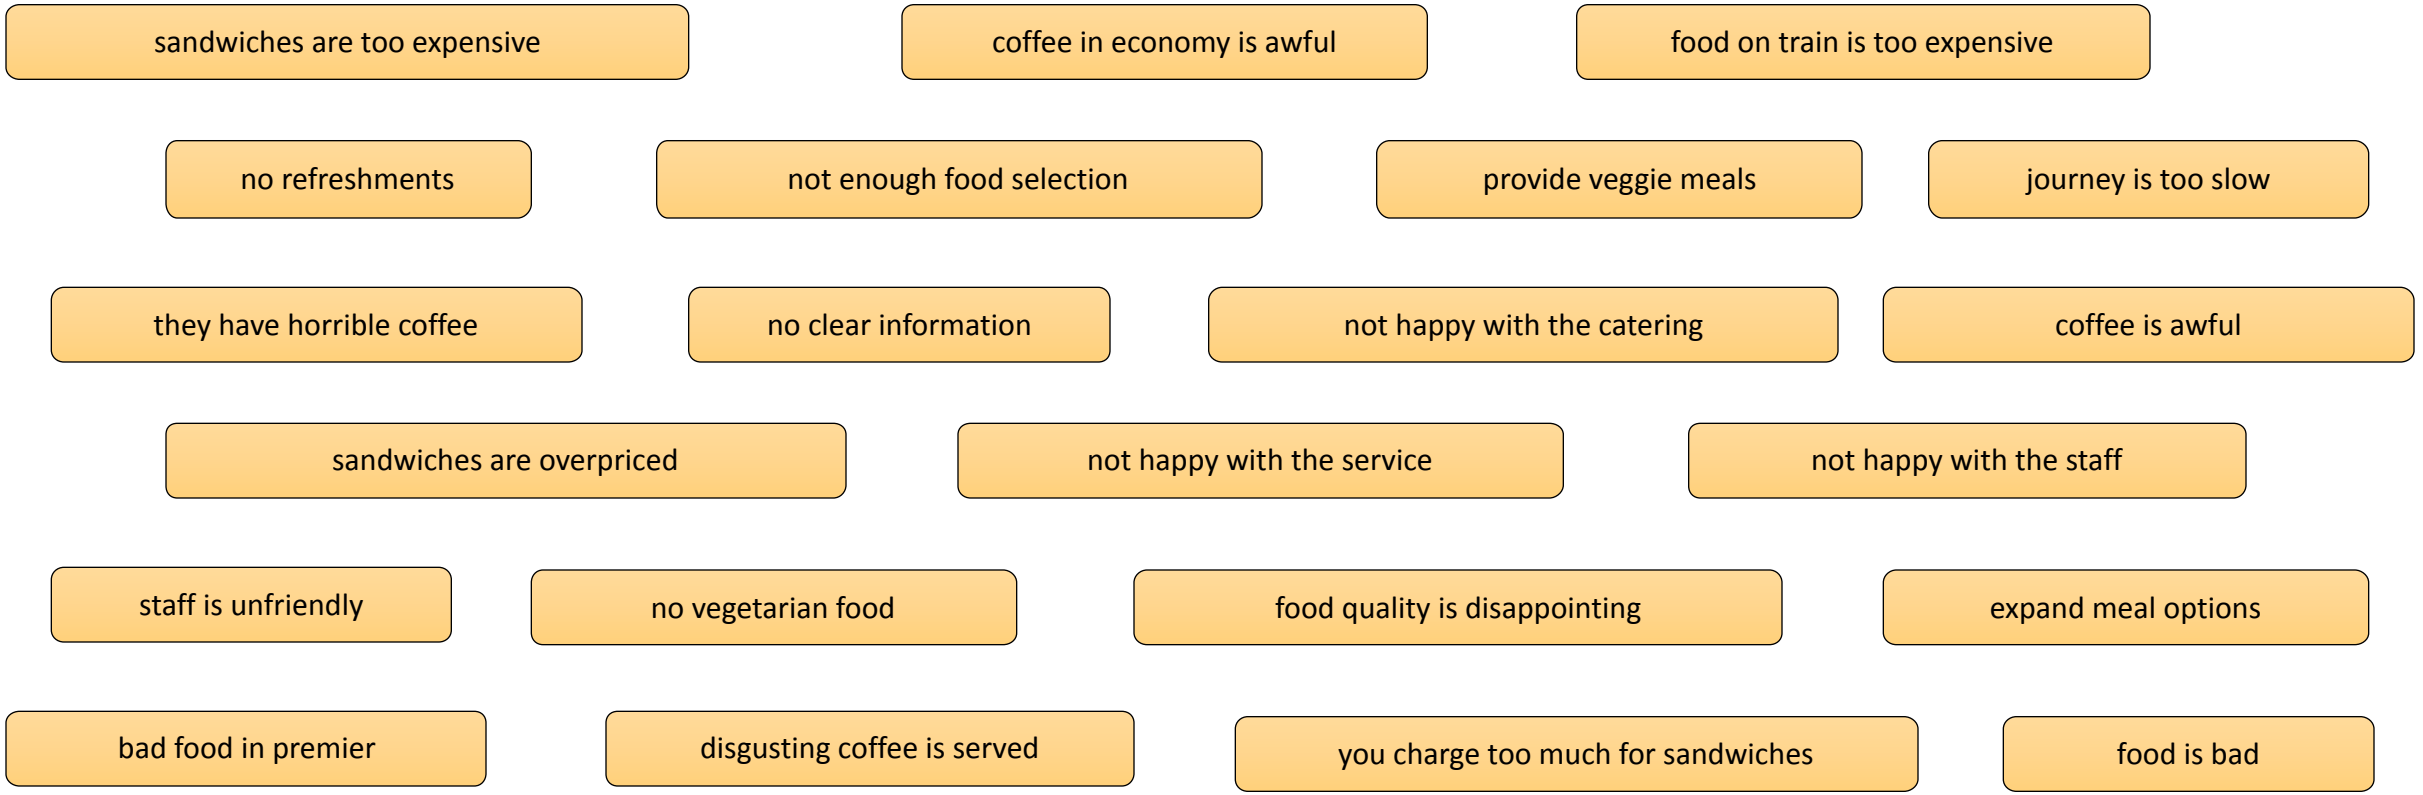

Novel Application: Text Exploration

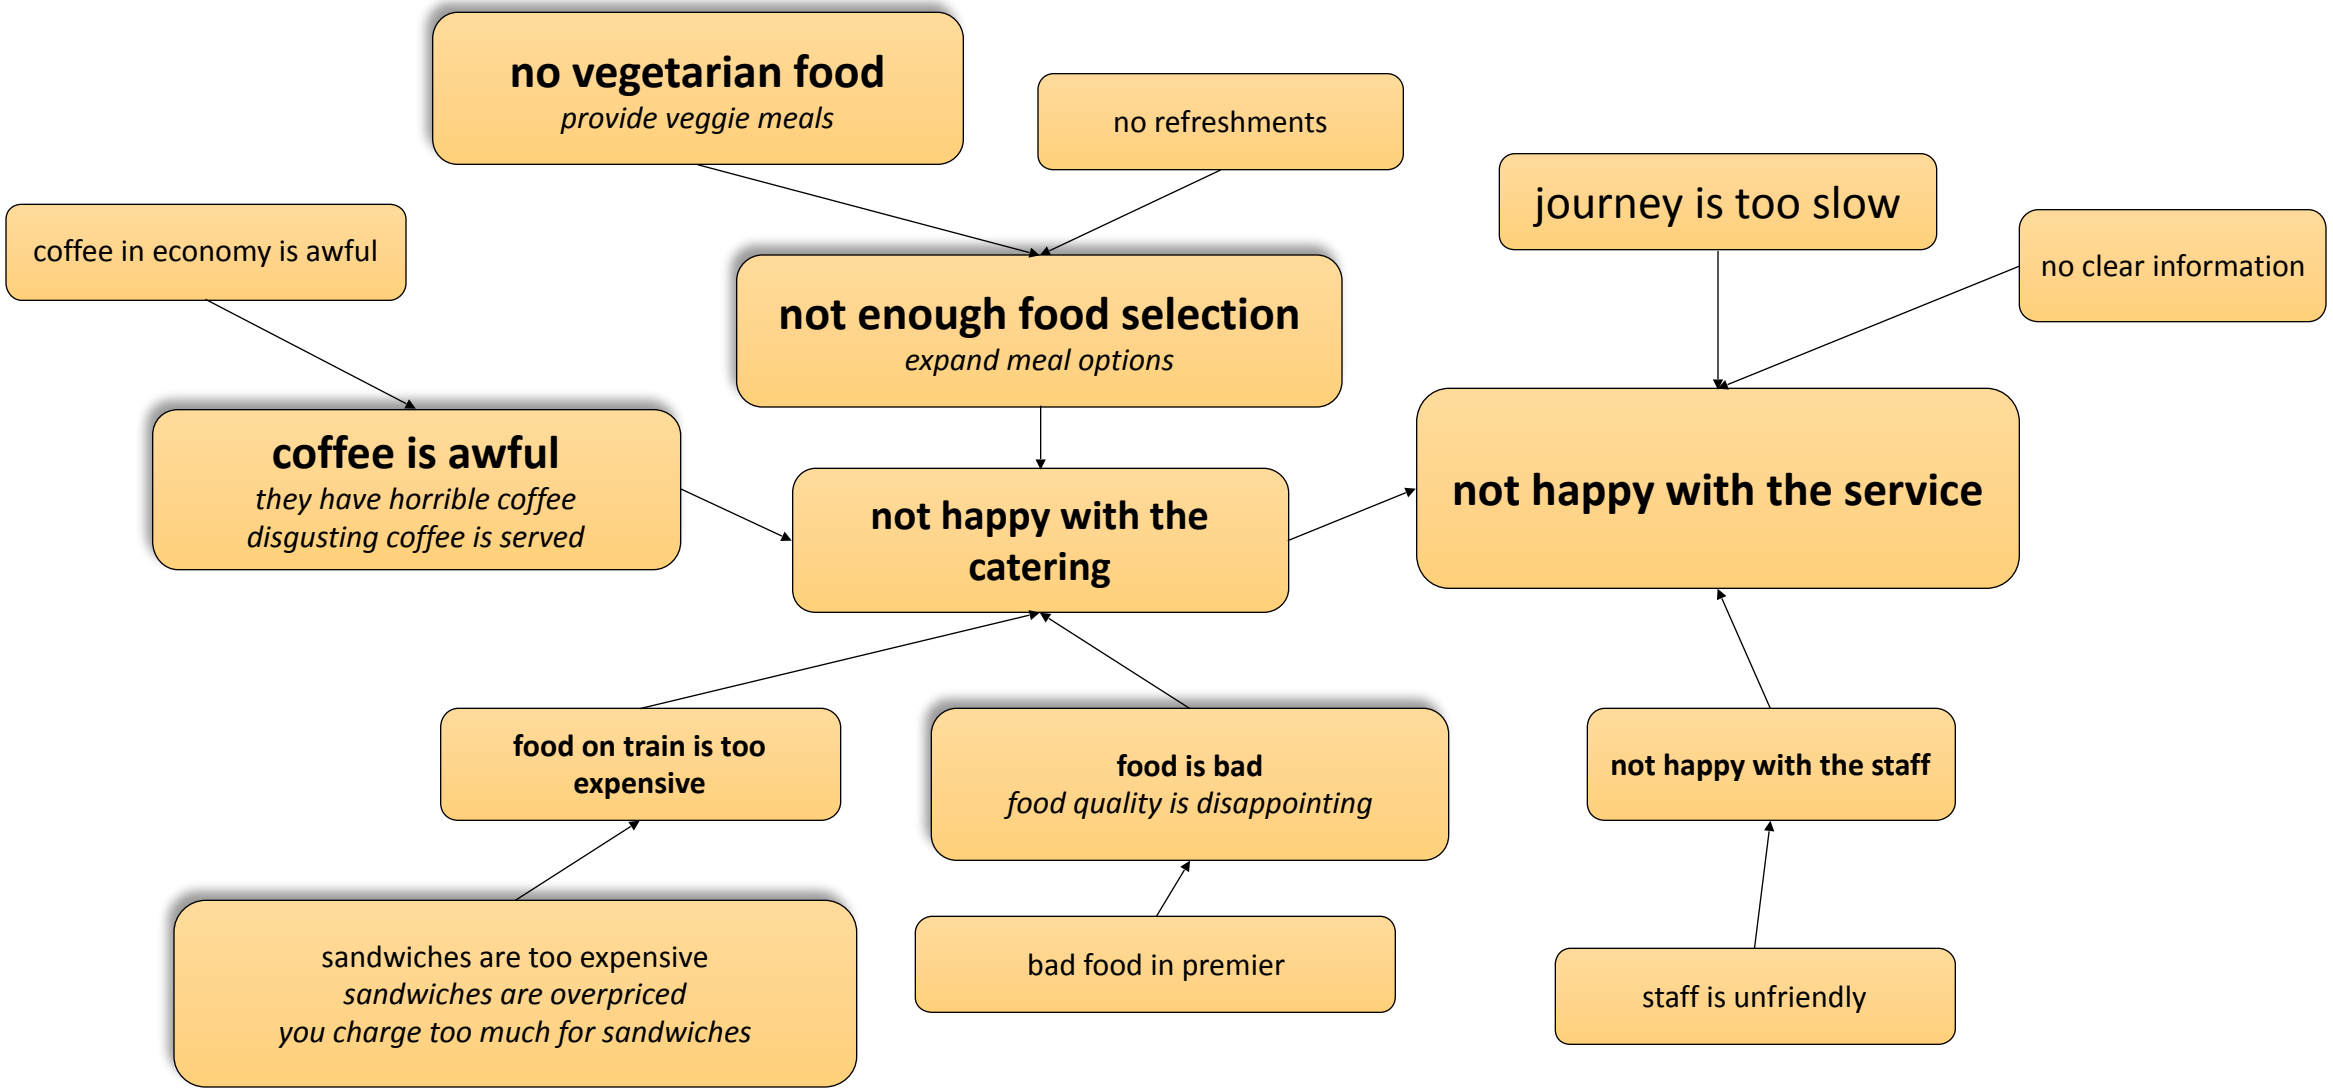

The EXCITEMENT Project

The EXCITEMENT Project

- Scientific goals
  - Advance textual entailment research
  - **Provide a flexible open platform for textual inference (EOP)**

The EXCITEMENT Project

- Scientific goals
  - Advance textual entailment research
  - **Provide a flexible open platform for textual inference (EOP)**
- Industrial goals
  - Advance customer interaction analytics, via
  - textual inference technologies

The EXCITEMENT Project

- Scientific goals
  - Advance textual entailment research
  - **Provide a flexible open platform for textual inference (EOP)**
- Industrial goals
  - Advance customer interaction analytics, via
  - textual inference technologies

The EXCITEMENT Project

- Scientific goals
  - Advance textual entailment research
  - **Provide a flexible open platform for textual inference (EOP)**
- Industrial goals
  - Advance customer interaction analytics, via
  - textual inference technologies

EXCITEMENT:
EXploring Customer Interactions via TExtual entailMENT

Outline

- Entailment recognition algorithm
  - Alignment based
- Entailment knowledge resources
- The EXCITEMENT Open Platform (EOP)
- Entailment graphs

Alignment-based Entailment Recognition

Alignment-based Entailment

- Various algorithms proposed to recognize textual entailment
- Recent work in EXCITEMENT: Alignment-based entailment
- Intuition: The more material in the hypothesis can be “explained” / ”covered” by the premise, the more likely entailment is

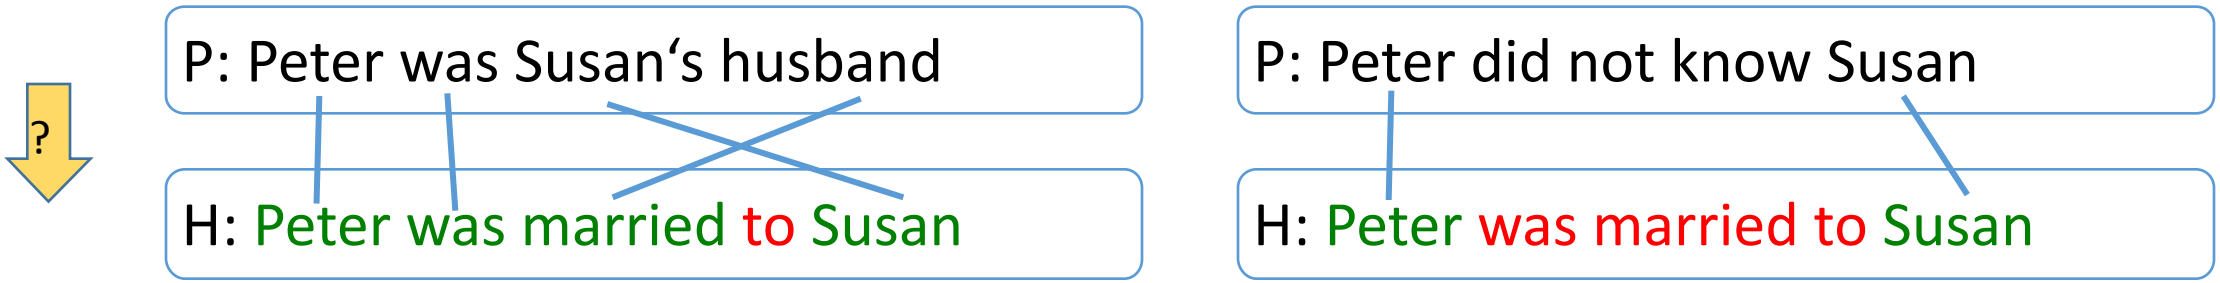

Alignment-based Entailment: The Algorithmic Level

- **Step 1**: Automatic linguistic analysis (Optional)
  - Normalize surface forms, detect structure

Part-of-speech tagger

Lemmatizer

Parser

...

NE V NE NN
P: Peter was Susan‘s husband

H: Peter was married to Susan
NE V V P NE

Alignment-based Entailment: The Algorithmic Level

- **Step 2**: Identify links between words or phrases across the two texts
  - What words/phrases of P can explain words/phrases of H?

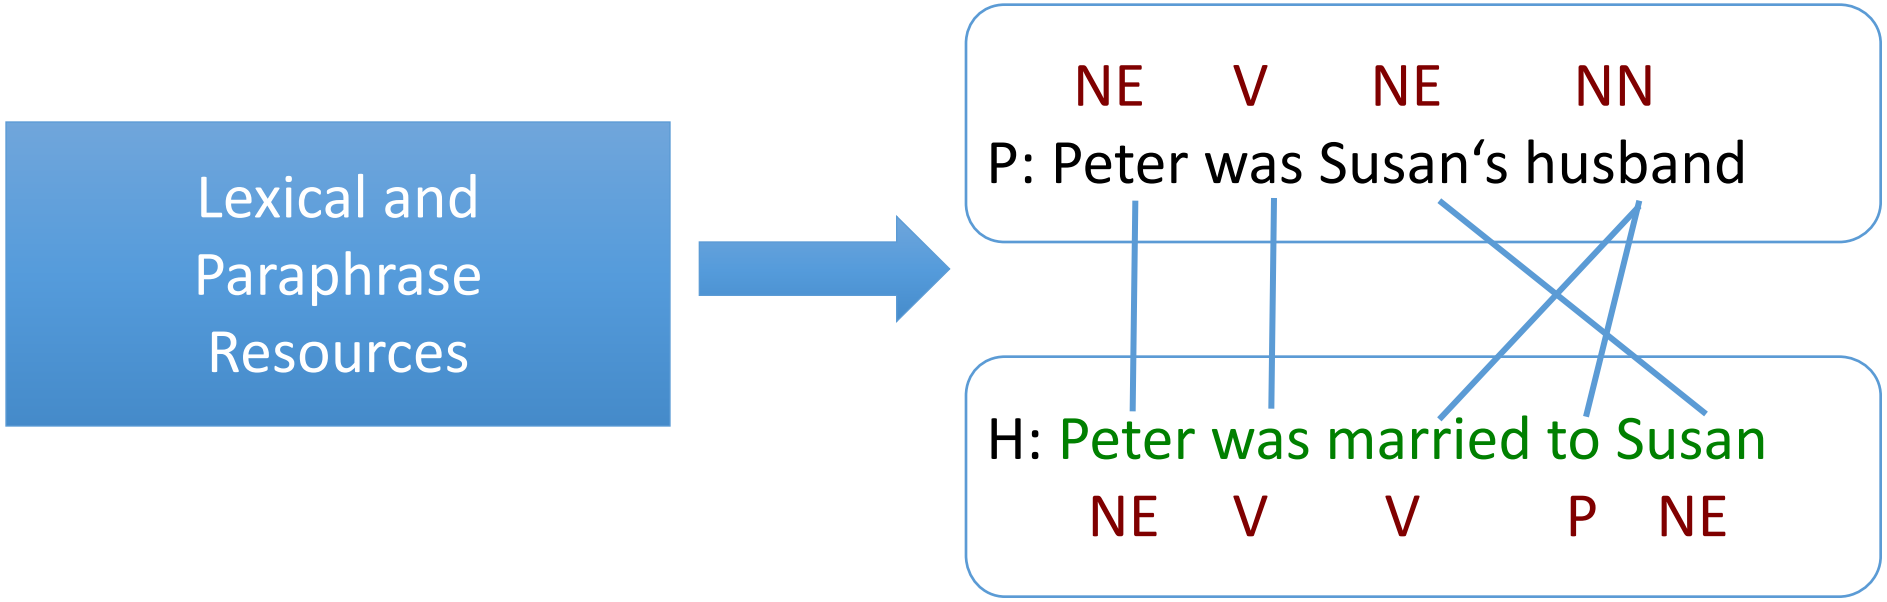

Lexical and Paraphrase Alignment Resources

- Broad-coverage knowledge needed to align words/phrases
  - Align identical words
  - Align lexically related **words**: use lexical resources (WordNet, distributional similarity)
  - Align equivalent/related **phrases**: use paraphrase resources

Peter → Peter

dog → mammal
Paris → France

was → used to
husband → married to

Alignment-based Entailment: The Algorithmic Level

- **Step 3**: Computation of features over alignment
  - Formulate features that capture typical properties of valid entailments

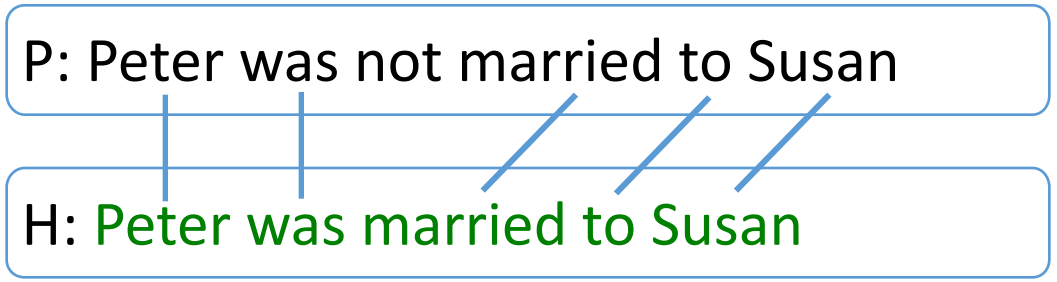

Concrete features

- Current implementation uses just four simple features

- **Word coverage**: What % of hypothesis words is covered?
- **Content word coverage**: What % of content words (N,V, A) covered?
- **Verb coverage**: What % of verbs is covered?
  - Verbs express the relations
- **Proper Noun coverage**: What % of proper nouns is covered?
  - Proper nouns express participants, typically require explicit mentions

- More features under development
  - E.g compatibility of negations

Alignment-based Entailment: The Algorithmic Level

- **Step 3**: Computation of features over alignment

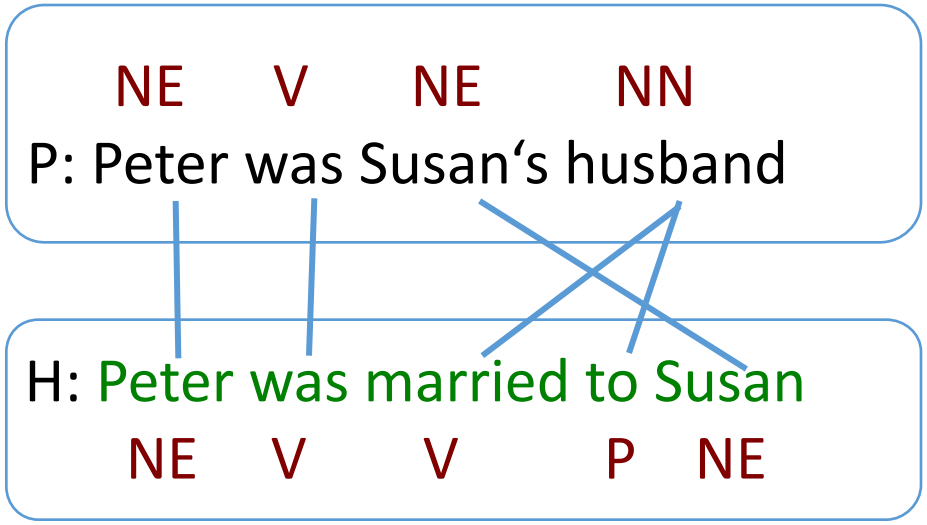

Word Coverage: 5/5 = 100%
Content Word Coverage: 4/4 =100%
Verb Coverage: 1/1=100%
Proper Noun Coverage: 2/2=100%

Alignment-based Entailment: The Algorithmic Level

- **Step 4**: Classification (logistic regression, with training examples)

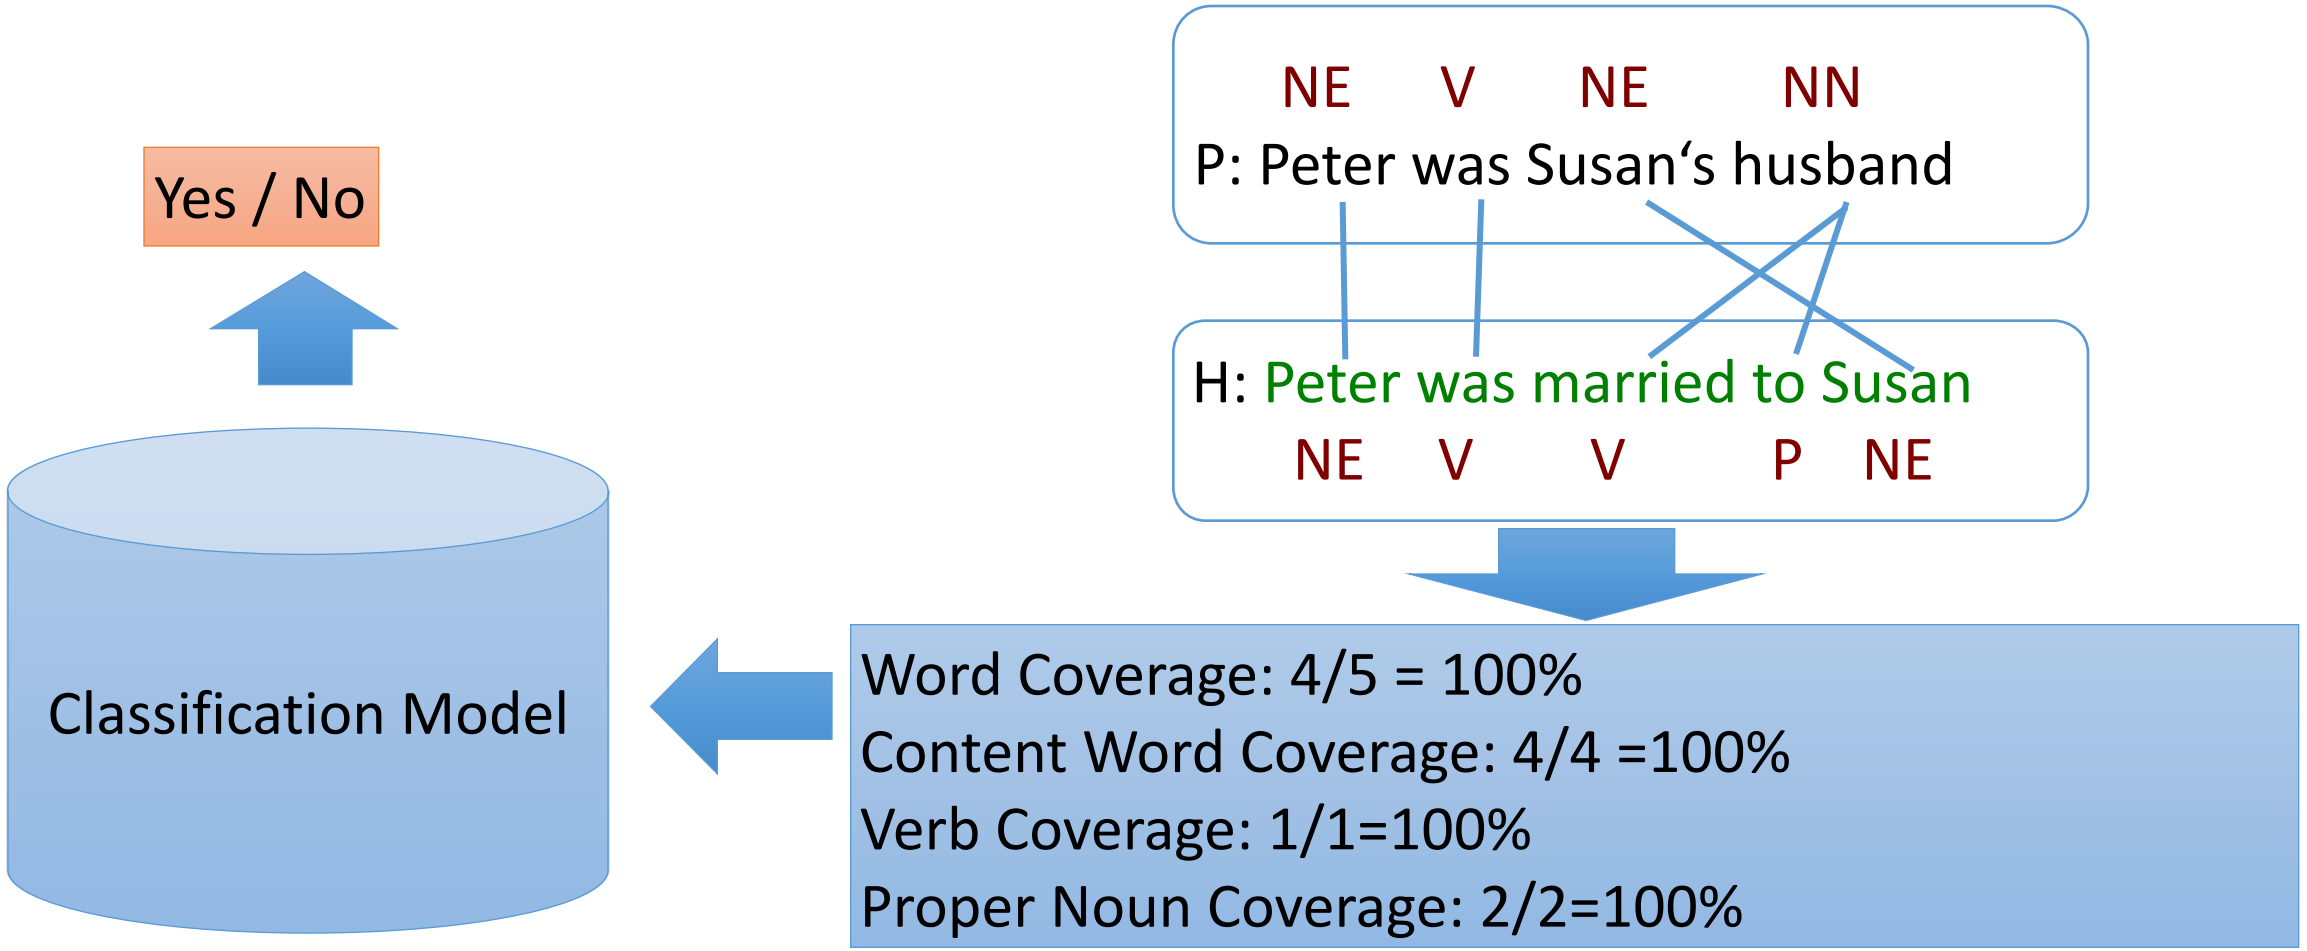

Why Alignment-based Entailment Recognition?

- Efficient
- (Almost completely) language-agnostic
- Robust: Can deal with noisy input data
  - Shallow linguistic cues
- Adaptable to new domains
  - Encode domain knowledge as alignment resource
- **Extensible**
- **State of the art useful accuracy**

- Will be included in EOP release in December 2014

Why Alignment-based Entailment Recognition?

- Efficient
- (Almost completely) language-agnostic
- Robust: Can deal with noisy input data
  - Shallow linguistic cues
- Adaptable to new domains
  - Encode domain knowledge as alignment resource
- **Extensible**
- **State of the art useful accuracy**

- Will be included in EOP release in December 2014

Why Alignment-based Entailment Recognition?

- Efficient
- (Almost completely) language-agnostic
- Robust: Can deal with noisy input data
  - Shallow linguistic cues
- Adaptable to new domains
  - Encode domain knowledge as alignment resource
- **Extensible**
- **State of the art useful accuracy**

- Will be included in EOP release in December 2014

HAS BEEN INCLUDED IN EOP 1.2 !

Extensibility

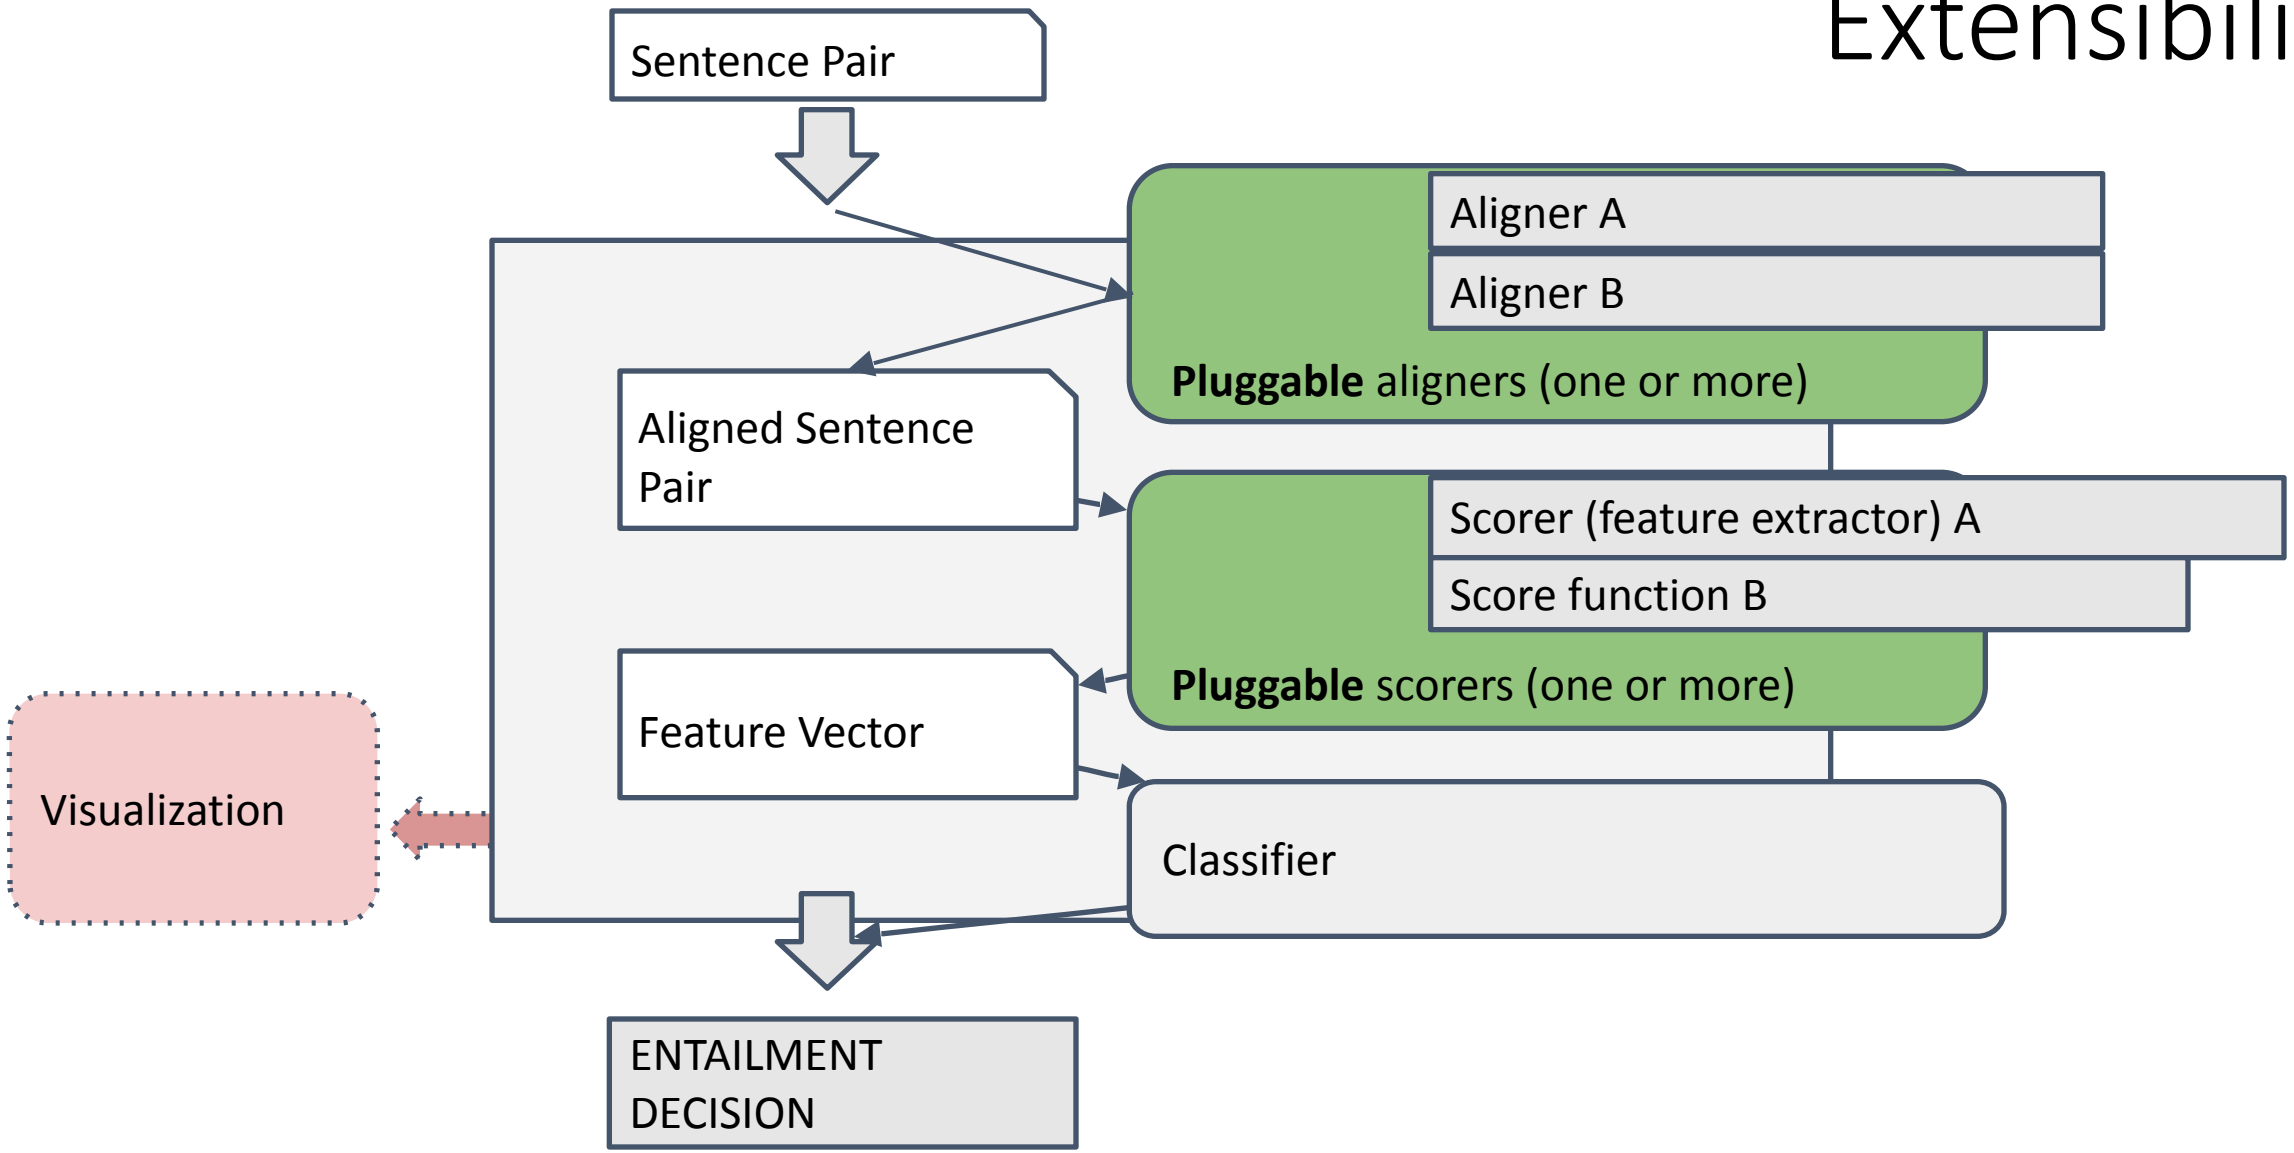

Performance at state-of-the-art [Dataset: RTE-3]

	Best Alignment-based EDA settings	Best previous EOP result
EN	**67.0**	66.8 (BIUTEE transformation)
IT	**65.4**	63.5% (EDITS transformation)
DE	**63.9**	63.5 (TIE matching features)

Performance at state-of-the-art [Dataset: RTE-3]

	Best Alignment-based EDA settings	Best previous EOP result
EN	**67.0**	66.8 (BIUTEE transformation)
IT	**65.4**	63.5% (EDITS transformation)
DE	**63.9**	63.5 (TIE matching features)

- Used for entailment graph construction on customer interactions data

Performance at state-of-the-art [Dataset: RTE-3]

	Best Alignment-based EDA settings	Best previous EOP result
EN	**67.0**	66.8 (BIUTEE transformation)
IT	**65.4**	63.5% (EDITS transformation)
DE	**63.9**	63.5 (TIE matching features)

- Used for entailment graph construction on customer interactions data
- Results seem useful

Entailment Knowledge Resources

Various Resources Types

- Wordnet
  - *pepper → spice stock → share*
- Derivational morphology
  - *allergenic→ allergy acquire → acquisition*
- Corpus-based distributional similarity
  - As seen in tutorial
  - Similar to word2vec type of output; limited correlation with entailment/equivalence
  - *Directional* similarity, usually somewhat better
- Wikipedia derived
  - *Madonna → singer*
- Paraphrasing – bilingual based

Tools for constructing knowledge resources for domain corpora and languages

Extraction from Wikipedia

(Shnarch et al., 2009)

E.T. the Extra-Terrestrial

From Wikipedia, the free encyclopedia

(Redirected from E.T. (film))

E.T. the Extra-Terrestrial is a 1982 science fiction film co-produced and directed by Steven Spielberg, written by Melissa Mathison and starring Henry Thomas, Robert MacNaughton, Drew Barrymore, Dee Wallace and Peter Coyote. It tells the story of Elliott (played by Thomas), a lonely boy who befriends a friendly alien dubbed "E.T.", who is stranded on Earth. Elliott and hi

Extraterrestrial life

- ***Be-complement***
- ***Redirect***
- ***Parenthesis***
- ***Link***

Bilingual-based Paraphrases

- Intuition: p and p’ are paraphrases if both translate into same phrase t (a “pivot”)

- Procedure:
  1. Word- and phrase-align parallel corpus (e.g. English-German)
  2. Extract bilingual translation table
  3. Hop from English to German and back to obtain paraphrase table (plus probability)

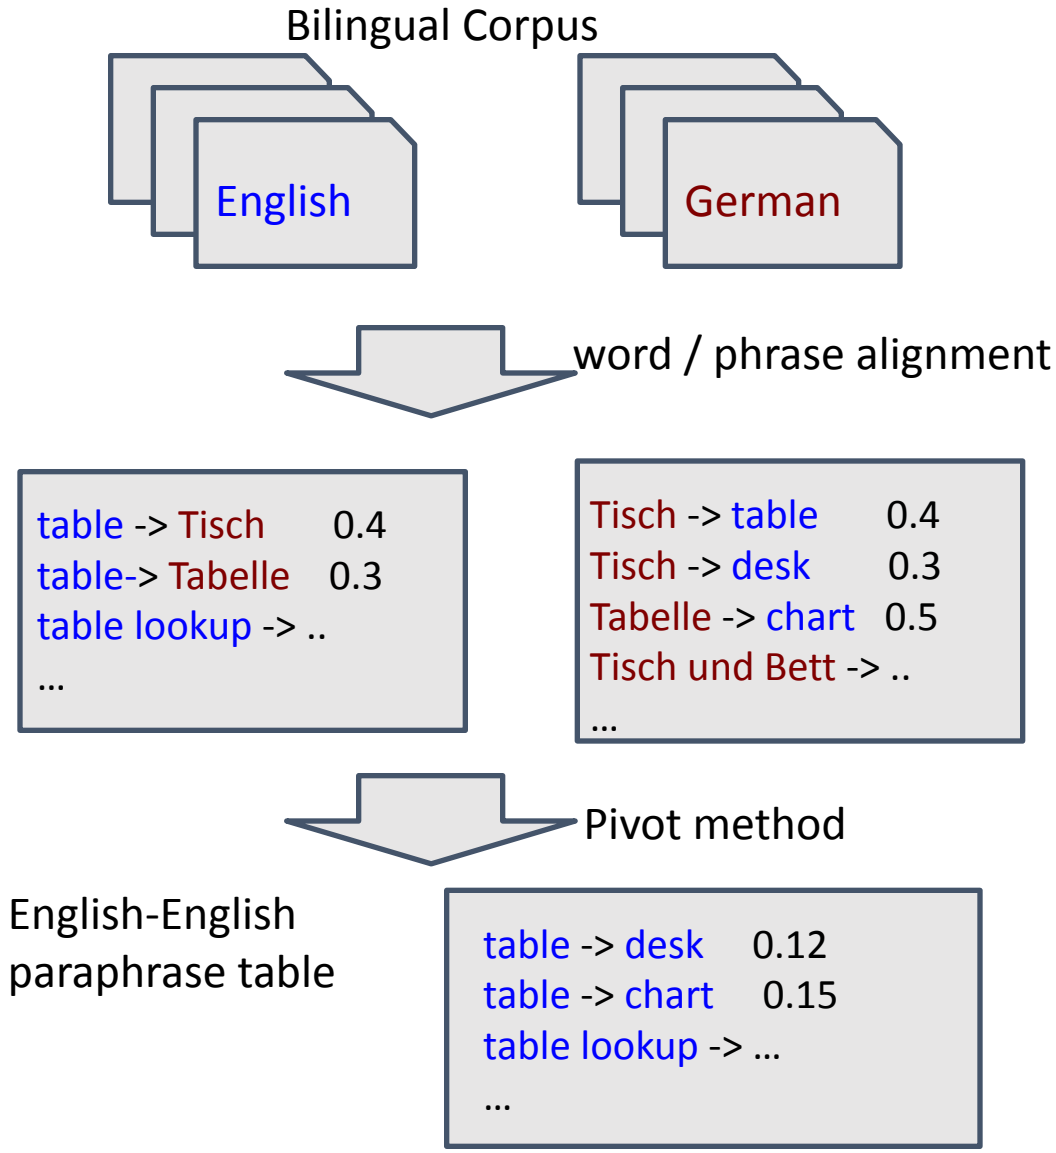

Transformation-based Approach

Recognizing Entailment via Sequences of Transformations

- Sequence of transformations

- Tree-Edits
  - Complete proofs
  - Estimate confidence
- Knowledge based Entailment Rules
  - Linguistically motivated
  - Formalize many types of knowledge

Recognizing Entailment via Sequences of Transformations

- Sequence of transformations

$$T = T_0 \rightarrow T_1 \rightarrow T_2 \rightarrow \ldots \rightarrow T_n = H$$

- Tree-Edits
  - Complete proofs
  - Estimate confidence
- Knowledge based Entailment Rules
  - Linguistically motivated
  - Formalize many types of knowledge

Example

Text: The boy was located by the police.

Hypothesis: Eventually, the police found the child.

Example

Text: The boy was located by the police.

The police located the boy.

The police found the boy.

The police found the child.

Hypothesis: Eventually, the police found the child.

Example

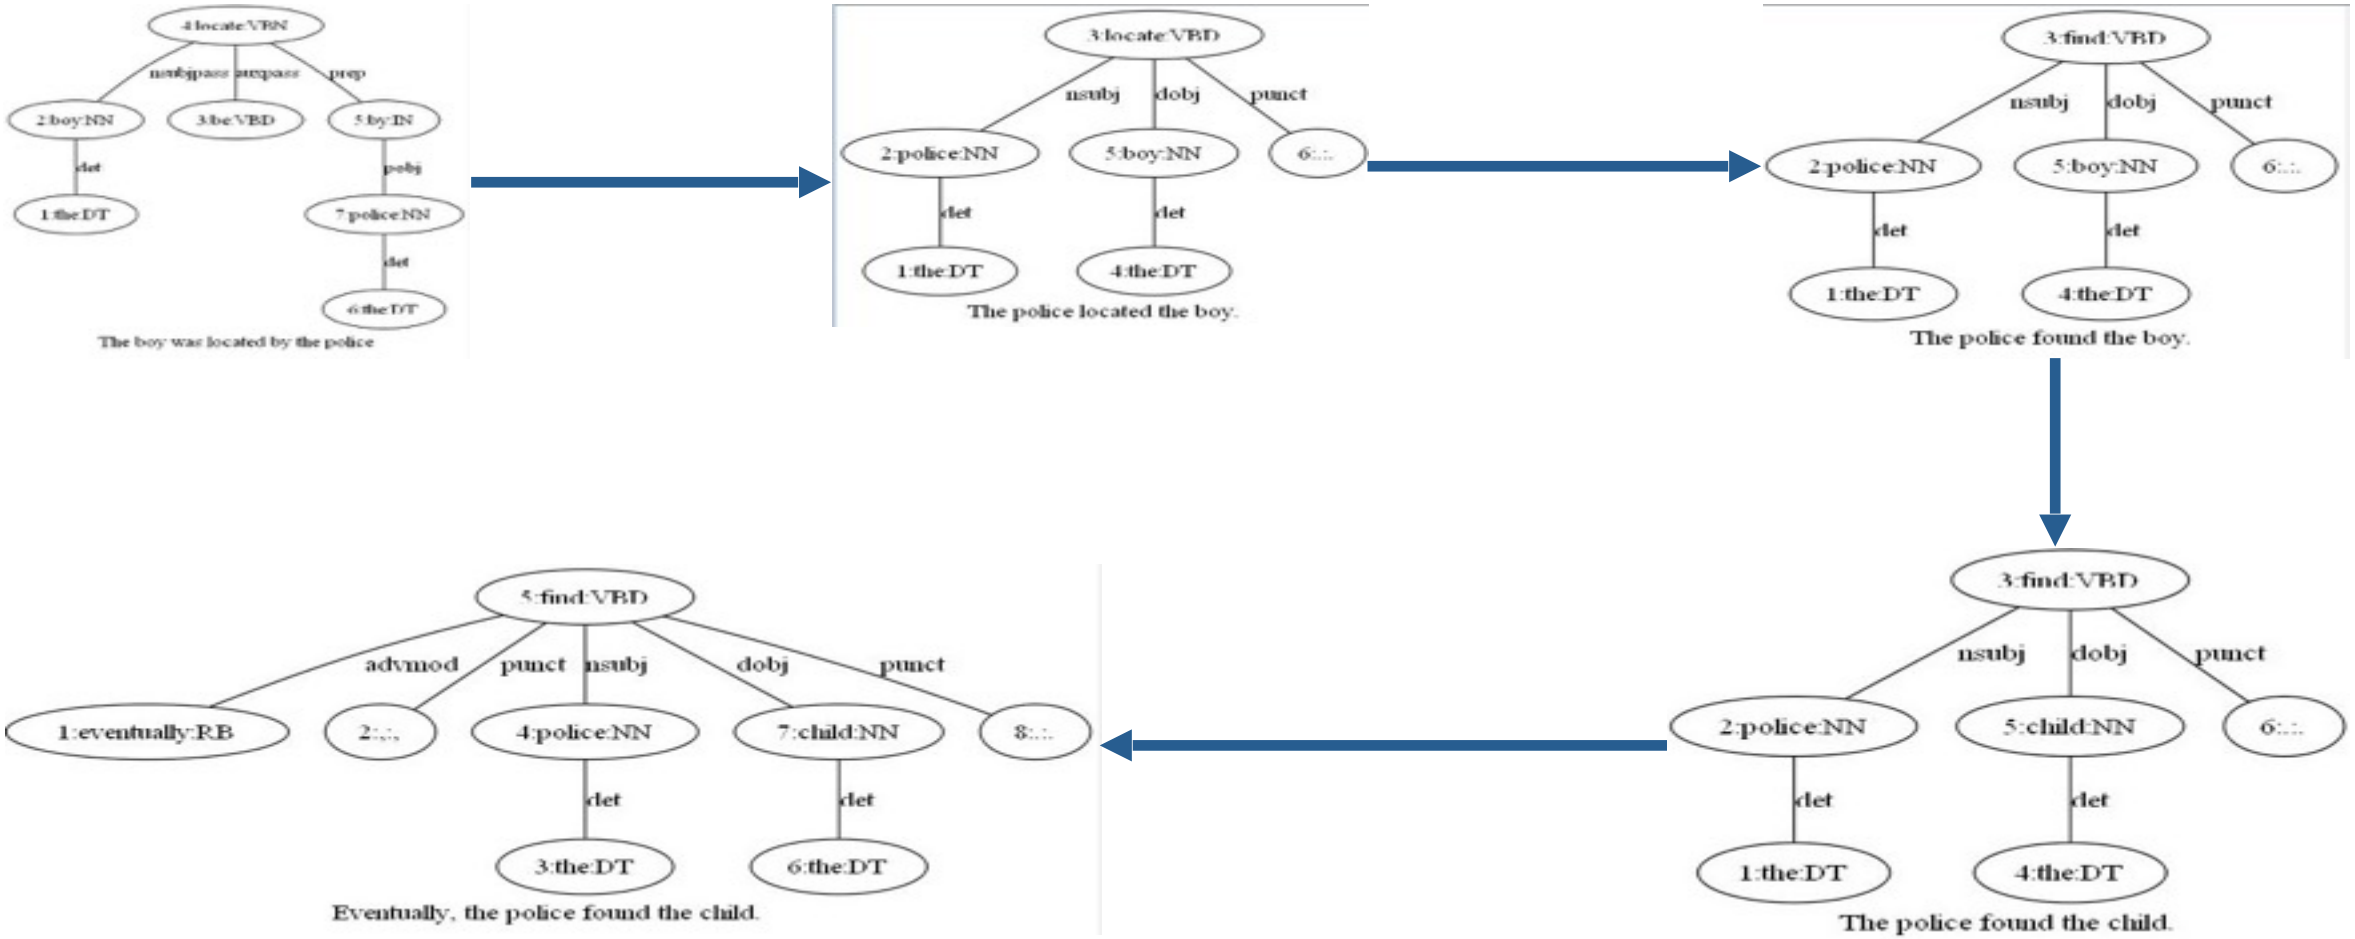
4:locate:VBN
nsubjpass
auxpass
prep
2:boy:NN
3:be:VBD
5:by:IN
det
pobj
1:the:DT
7:police:NN
det
6:the:DT
The boy was located by the police
3:locate:VBD
nsubj
dobj
punct
2:police:NN
5:boy:NN
6:.:.
det
det
1:the:DT
4:the:DT
The police located the boy.
3:find:VBD
nsubj
dobj
punct
2:police:NN
5:boy:NN
6:.:.
det
det
1:the:DT
4:the:DT
The police found the boy.
3:find:VBD
nsubj
dobj
punct
2:police:NN
5:child:NN
6:.:.
det
det
1:the:DT
4:the:DT
The police found the child.
5:find:VBD
advmod
punct
nsubj
dobj
punct
1:eventually:RB
2:,:,
4:police:NN
7:child:NN
8:.:.
det
det
3:the:DT
6:the:DT
Eventually, the police found the child.

Entailment Rules

Generic Syntactic

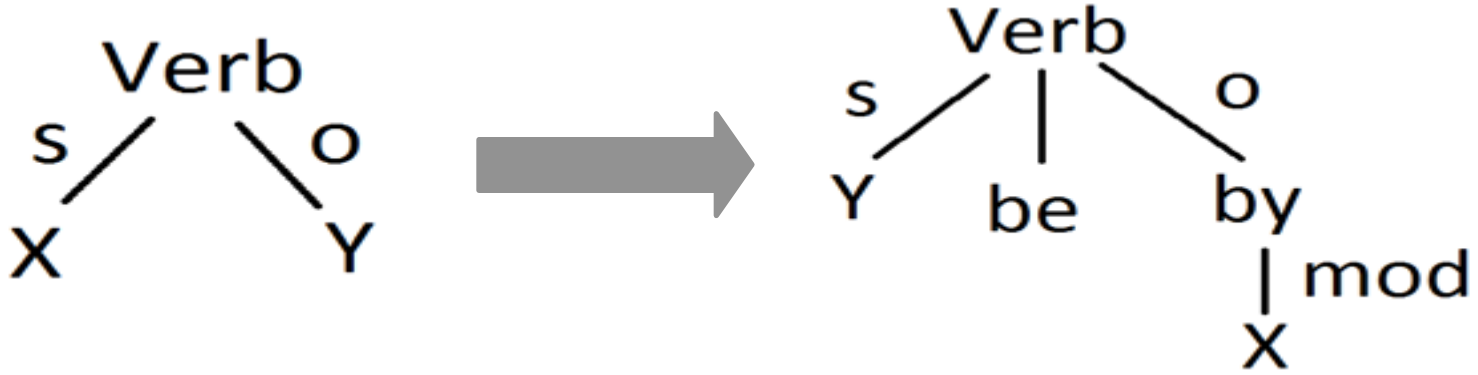

Lexical Syntactic

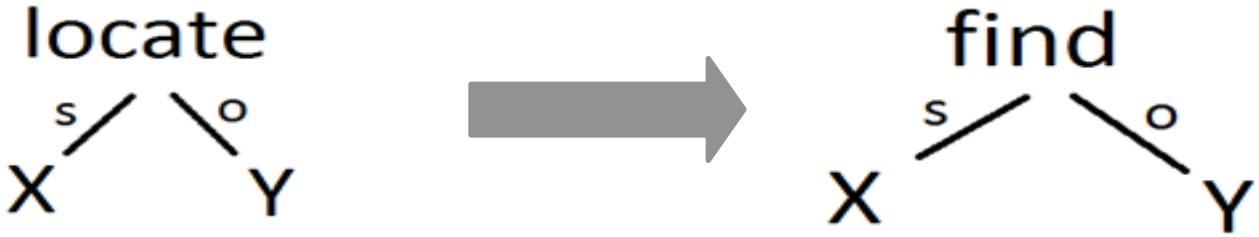

Lexical boy → child

Bar-Haim et al. 2007. ***Semantic inference at the lexical-syntactic level***.

Example

Text: The boy was located by the police.

Example

Text: The boy was located by the police.

Passive to active

The police located the boy.

X locate Y → X find Y

The police found the boy.

Boy → child

The police found the child.

Insertion on the fly

Example

Text: The boy was located by the police.

Passive to active

The police located the boy.

X locate Y → X find Y

The police found the boy.

Boy → child

The police found the child.

Insertion on the fly

Hypothesis: Eventually, the police found the child.

Recognizing Entailment as a Scientific Competition

Recognizing Textual Entailment (RTE) Challenge – A Scientific Competition

- Since 2005 until today - RTE-1 to RTE-7
- Main motivation: Bring together scientists from all over the world, in order to commonly push forward the scientific field of „applied semantics“ („open collaboration“).

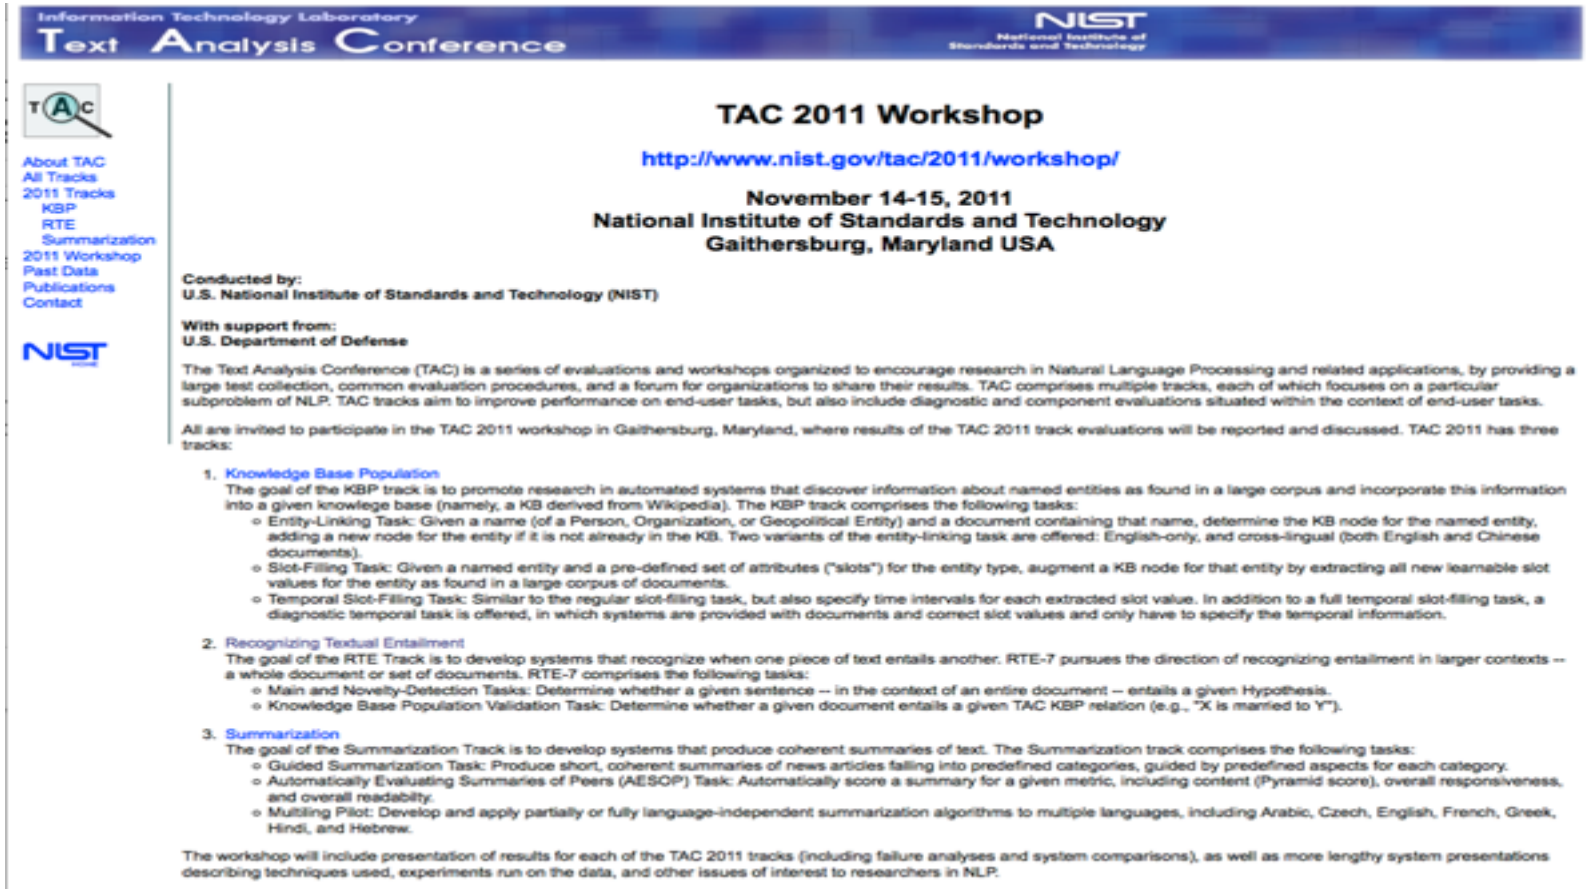

Information Technology Laboratory
Text Analysis Conference
NIST National Institute of Standards and Technology

About TAC
All Tracks
2011 Tracks
KBP
RTE
Summarization
2011 Workshop
Past Data
Publications
Contact

TAC 2011 Workshop

http://www.nist.gov/tac/2011/workshop/

November 14-15, 2011
National Institute of Standards and Technology
Gaithersburg, Maryland USA

Conducted by:
U.S. National Institute of Standards and Technology (NIST)

With support from:
U.S. Department of Defense

The Text Analysis Conference (TAC) is a series of evaluations and workshops organized to encourage research in Natural Language Processing and related applications, by providing a large test collection, common evaluation procedures, and a forum for organizations to share their results. TAC comprises multiple tracks, each of which focuses on a particular subproblem of NLP. TAC tracks aim to improve performance on end-user tasks, but also include diagnostic and component evaluations situated within the context of end-user tasks.

All are invited to participate in the TAC 2011 workshop in Gaithersburg, Maryland, where results of the TAC 2011 track evaluations will be reported and discussed. TAC 2011 has three tracks:

1. Knowledge Base Population
   The goal of the KBP track is to promote research in automated systems that discover information about named entities as found in a large corpus and incorporate this information into a given knowlege base (namely, a KB derived from Wikipedia). The KBP track comprises the following tasks:
   - Entity-Linking Task: Given a name (of a Person, Organization, or Geopolitical Entity) and a document containing that name, determine the KB node for the named entity, adding a new node for the entity if it is not already in the KB. Two variants of the entity-linking task are offered: English-only, and cross-lingual (both English and Chinese documents).
   - Slot-Filling Task: Given a named entity and a pre-defined set of attributes ("slots") for the entity type, augment a KB node for that entity by extracting all new learnable slot values for the entity as found in a large corpus of documents.
   - Temporal Slot-Filling Task: Similar to the regular slot-filling task, but also specify time intervals for each extracted slot value. In addition to a full temporal slot-filling task, a diagnostic temporal task is offered, in which systems are provided with documents and correct slot values and only have to specify the temporal information.
2. Recognizing Textual Entailment
   The goal of the RTE Track is to develop systems that recognize when one piece of text entails another. RTE-7 pursues the direction of recognizing entailment in larger contexts -- a whole document or set of documents. RTE-7 comprises the following tasks:
   - Main and Novelty-Detection Tasks: Determine whether a given sentence -- in the context of an entire document -- entails a given Hypothesis.
   - Knowledge Base Population Validation Task: Determine whether a given document entails a given TAC KBP relation (e.g., "X is married to Y").
3. Summarization
   The goal of the Summarization Track is to develop systems that produce coherent summaries of text. The Summarization track comprises the following tasks:
   - Guided Summarization Task: Produce short, coherent summaries of news articles falling into predefined categories, guided by predefined aspects for each category.
   - Automatically Evaluating Summaries of Peers (AESOP) Task: Automatically score a summary for a given metric, including content (Pyramid score), overall responsiveness, and overall readability.
   - Multiling Pilot: Develop and apply partially or fully language-independent summarization algorithms to multiple languages, including Arabic, Czech, English, French, Greek, Hindi, and Hebrew.

The workshop will include presentation of results for each of the TAC 2011 tracks (including failure analyses and system comparisons), as well as more lengthy system presentations describing techniques used, experiments run on the data, and other issues of interest to researchers in NLP.

Differences between RTE-1-5 and RTE-6-7

RTE1-5 vs. RTE6 Main Task

RTE1-5	RTE6
• RTE on isolated T-H pairs	• RTE within a corpus
• T-H pairs drawn from multiple applications	• Summarization application setting
• T and H do not contain references to information outside the pair itself	• Both T and H are to be interpreted within the context of the corpus
• The distribution of entailment is determined a priori	• Reflects the natural distribution of entailment in a corpus

NIST - November 16, 2010 RTE-6@TAC2010

Data format for RTE-1-5

```
<pair id="1" entailment="YES" task="IE" length="short" >
<t>The sale was made to pay Yukos' US$ 27.5 billion tax bill, Yuganskneftegaz was originally sold for
US$ 9.4 billion to a little known company Baikalfinansgroup which was later bought by the Russian
state-owned oil company Rosneft .</t>
<h>Baikalfinansgroup was sold to Rosneft.</h> </pair>

<pair id="2" entailment="NO" task="IE" length="short" >
<t>The sale was made to pay Yukos' US$ 27.5 billion tax bill, Yuganskneftegaz was originally sold for
US$9.4 billion to a little known company Baikalfinansgroup which was later bought by the Russian
state-owned oil company Rosneft .</t>
<h>Yuganskneftegaz cost US$ 27.5 billion.</h> </pair>

<pair id="3" entailment="NO" task="IE" length="long" >
<t>Loraine besides participating in Broadway's Dreamgirls, also participated in the Off-Broadway
production of "Does A Tiger Have A Necktie". In 1999, Loraine went to London, United Kingdom. There
she participated in the production of "RENT" where she was cast as "Mimi" the understudy.</t>
<h>"Does A Tiger Have A Necktie" was produced in London.</h> </pair>

<pair id="4" entailment="YES" task="IE" length="long" >
<t>"The Extra Girl" (1923) is a story of a small-town girl, Sue Graham (played by Mabel Normand) who
comes to Hollywood to be in the pictures. This Mabel Normand vehicle, produced by Mack Sennett,
followed earlier films about the film industry and also paved the way for later films about Hollywood,
such as King Vidor's "Show People" (1928).</t>
<h>"The Extra Girl" was produced by Sennett.</h> </pair>
```

RTE-6 Example

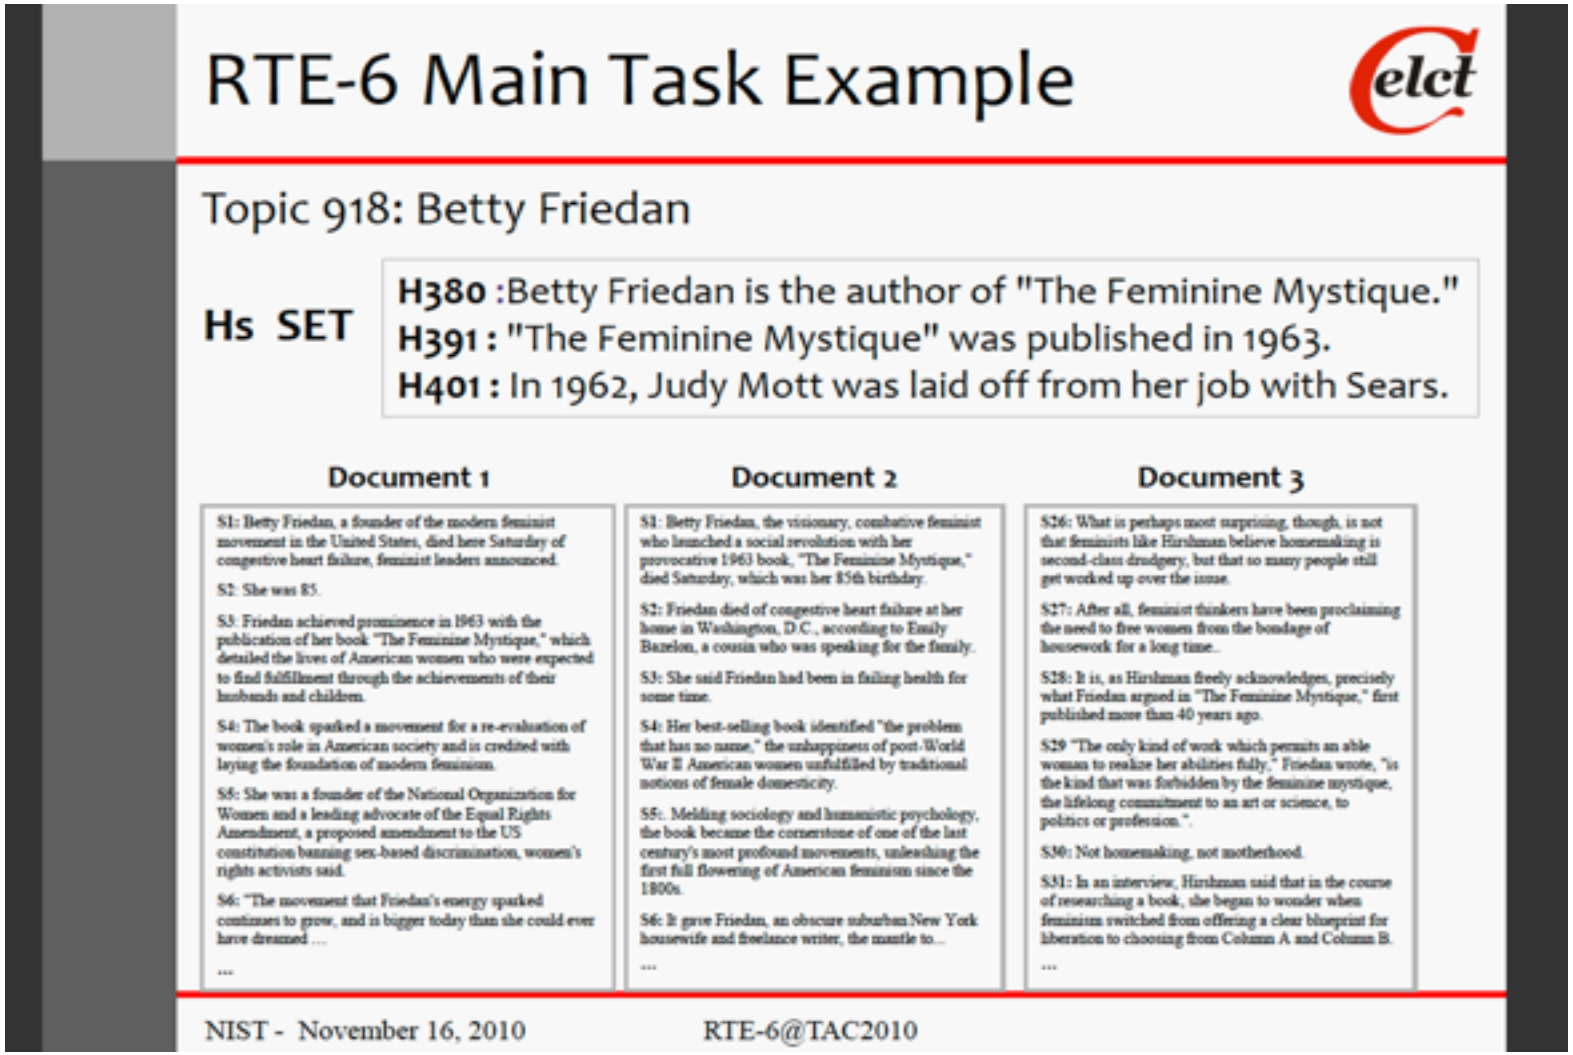

RTE-6 Example

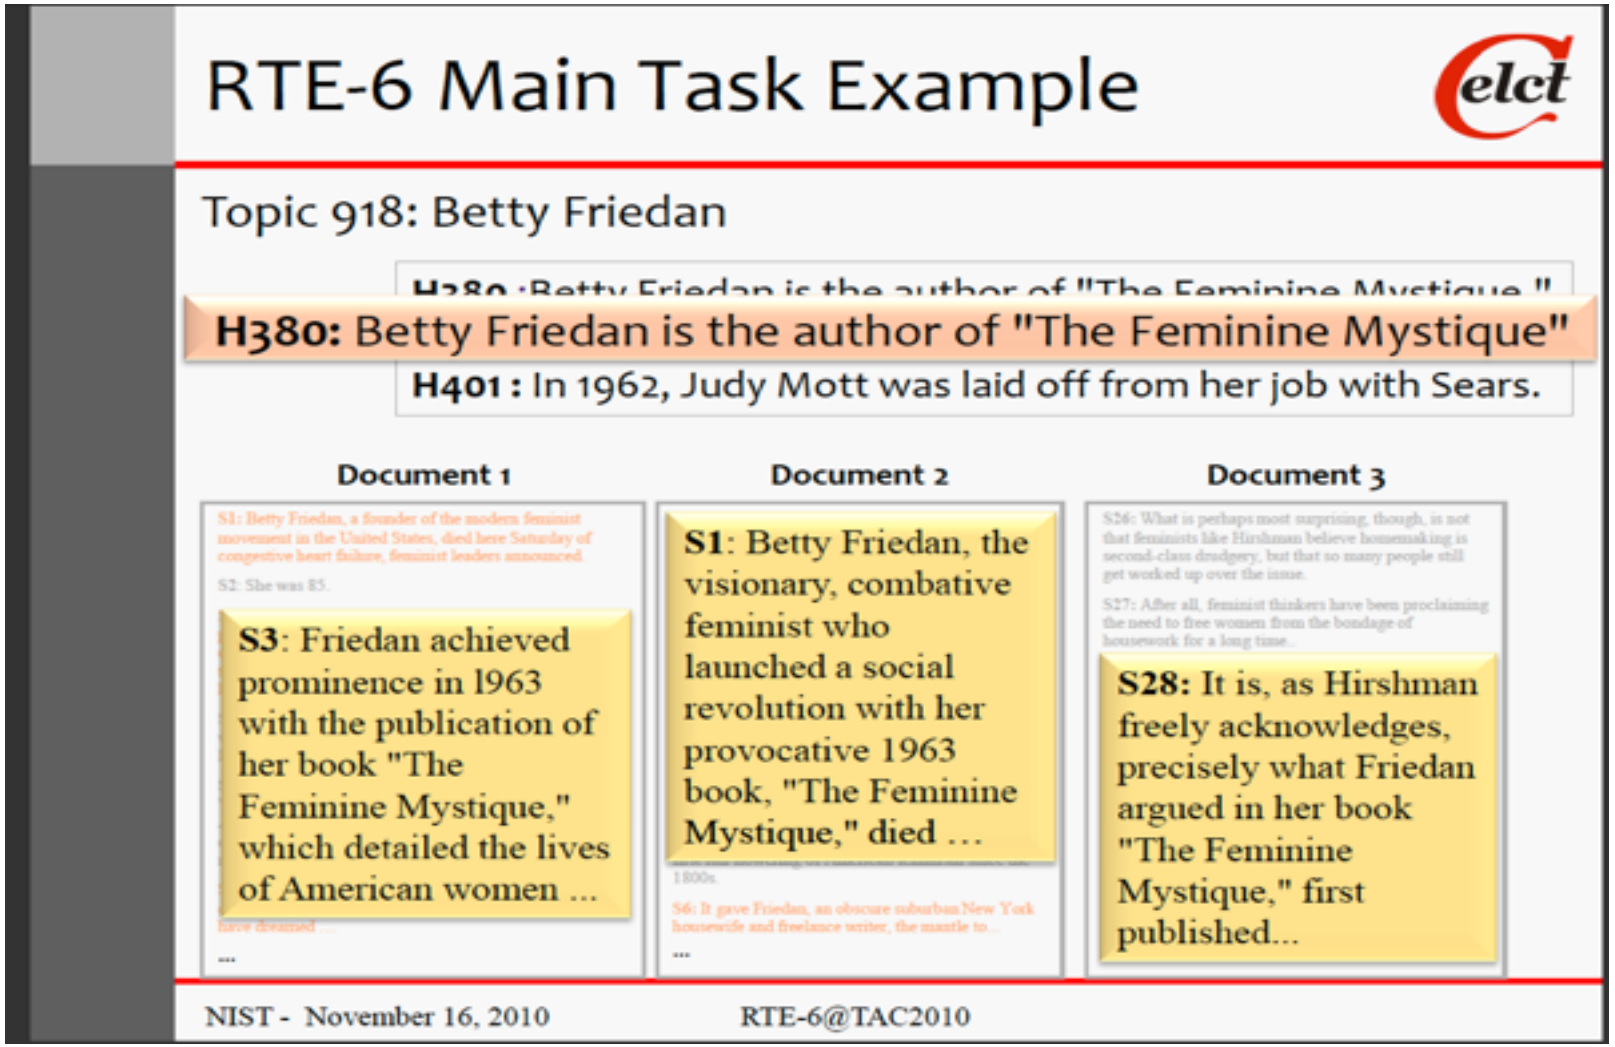

Example system - DFKI-LITE - RTE-7

- A single machine learning engine (a linear SVM) is fed with features extracted from many different sources and learns to select the best (Volokh & Neumann, 2011).
- **It also uses METEOR as text aligner!**

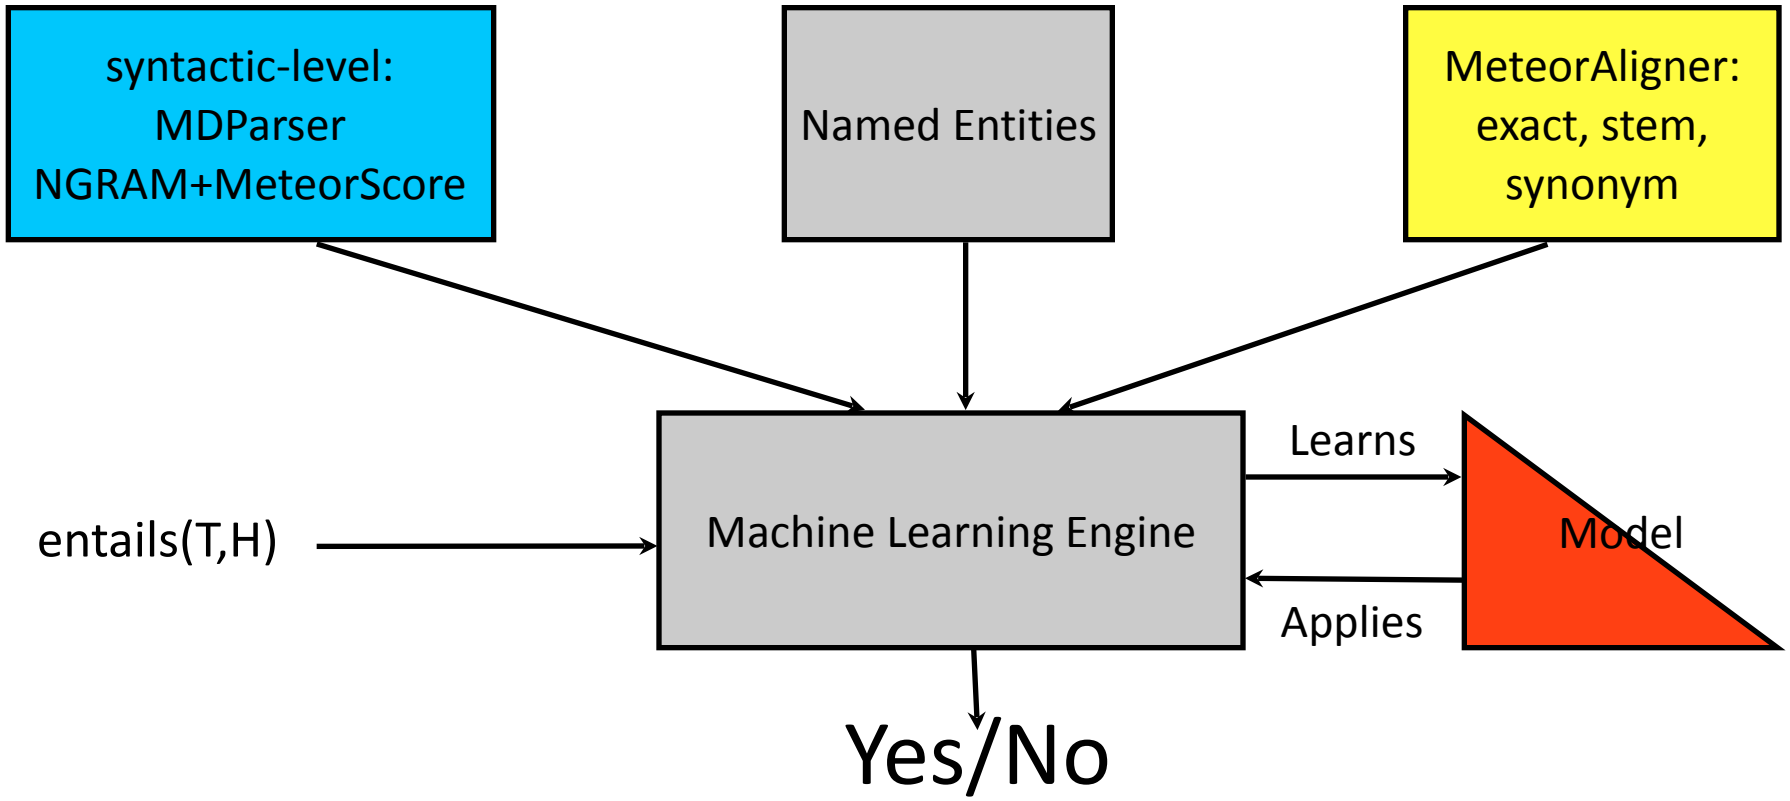

Main Task Evaluation

13 participants (33 runs)

- Evaluation measures:
  - Precision, Recall, F-measure (micro-averaged)
- IR Baselines:

	Precision	Recall	F1
Lucene_5	37.00	37.84	37.41
Lucene_10	27.07	55.20	36.33
Lucene_15	21.15	64.65	31.85
Lucene_20	17.71	71.64	28.40
Lucene_100	5.83	100	11.02

NIST - November 14, 2011 RTE-7@TAC2011

- 43.41 micro-average F1-score
- 46.34 macro-average F1-score

  → Above median, big improvement over the last year
- Very robust solution to an exremely large amount of data

  - >50% can be solved this way if account for weaknesses
- Problem-specific alternatives can still be included for the rest of the data

Best Results

Team	Precision	Recall	F-measure
IKOMA1	46.96	49.08	48.00
u_tokyo3	46.84	43.58	45.15
BUPTTeam1	45.02	44.95	44.99
CELI1	41.88	46.56	44.10
DFKI2	50.77	37.92	43.41
BIU2	41.81	44.11	42.93
FBK_irst3	46.59	38.07	41.90
Baseline_Lucene5	*30.78*	*39.58*	*34.63*
te_iitb1	20.67	60.24	30.78
JU_CSE_TAC2	26.66	35.55	30.47
ICL1	47.88	21.56	29.73
UAIC20112	30.21	25.84	27.85
SJTU_CIT3	17.92	33.33	23.31
SINAI3	47.3	8.72	14.72
Baseline_LuceneAll	*4.73*	*100.00*	*9.03*

NIST - November 14, 2011 RTE-7@TAC2011

DFKI-LITE obtained highest precision

NOTE:

The next slides are optional and will not be covered in the exam

Excitement Open Platform

Excitement Open Platform (EOP)

- **Excitement Project**: develop generic entailment platform
  - Step 1: Decouple preprocessing and actual entailment computation
  - Step 2: Decompose inference into components

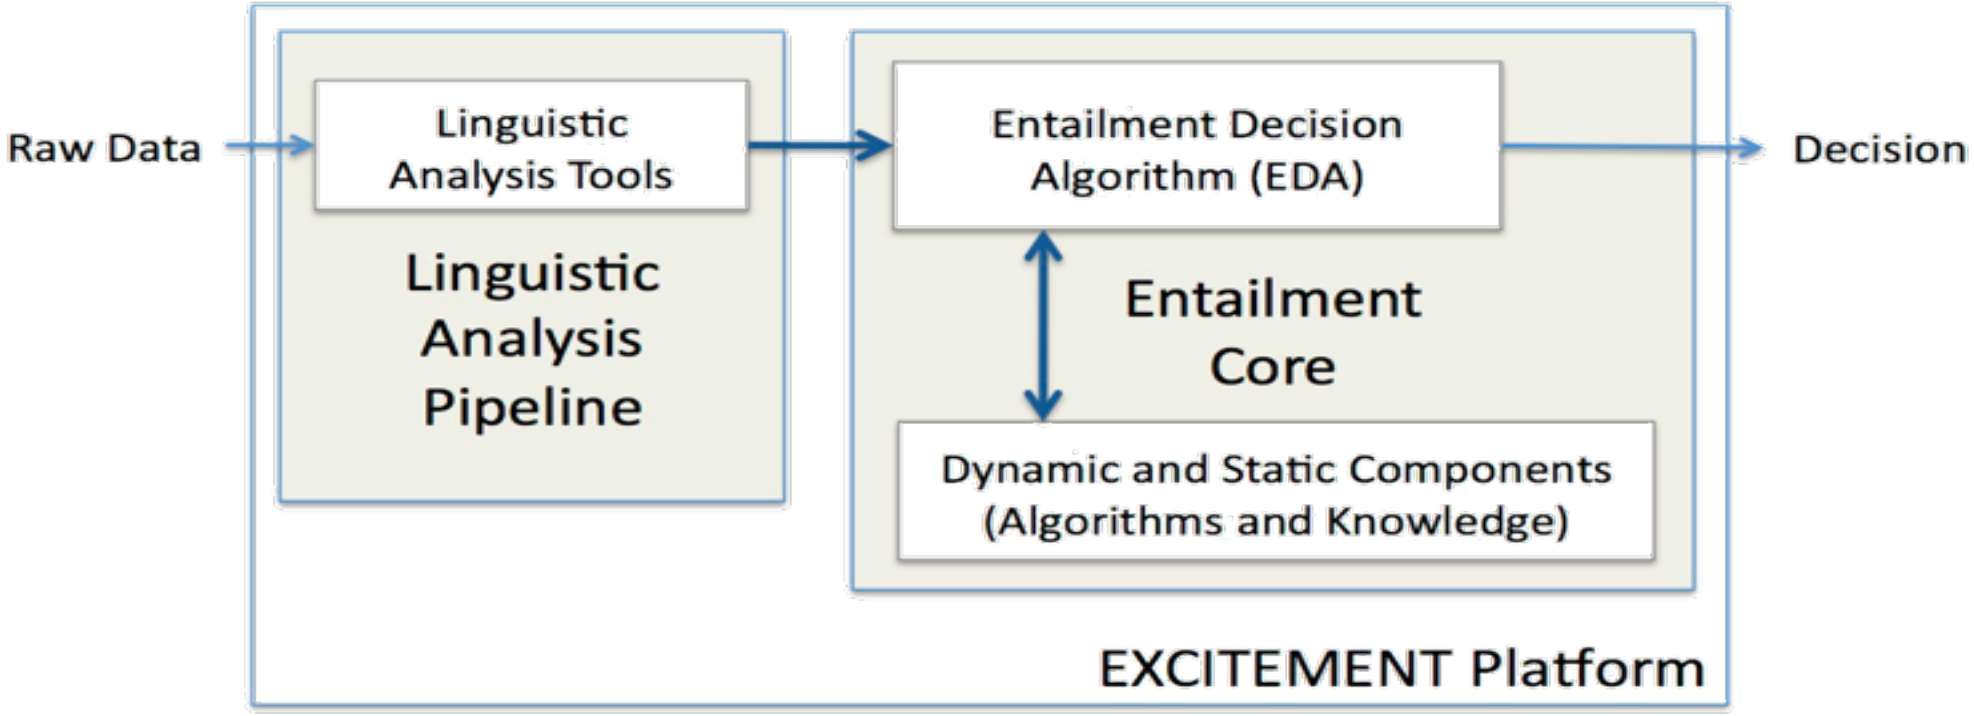

EXCITEMENT EU project: http://www.excitement-project.eu

Magnini et al.: The Excitement Open Platform, ACL demo 2014

Pado et al.: Journal Natural Language Engineering, 2014

EXCITEMENT Platform for Textual Inference

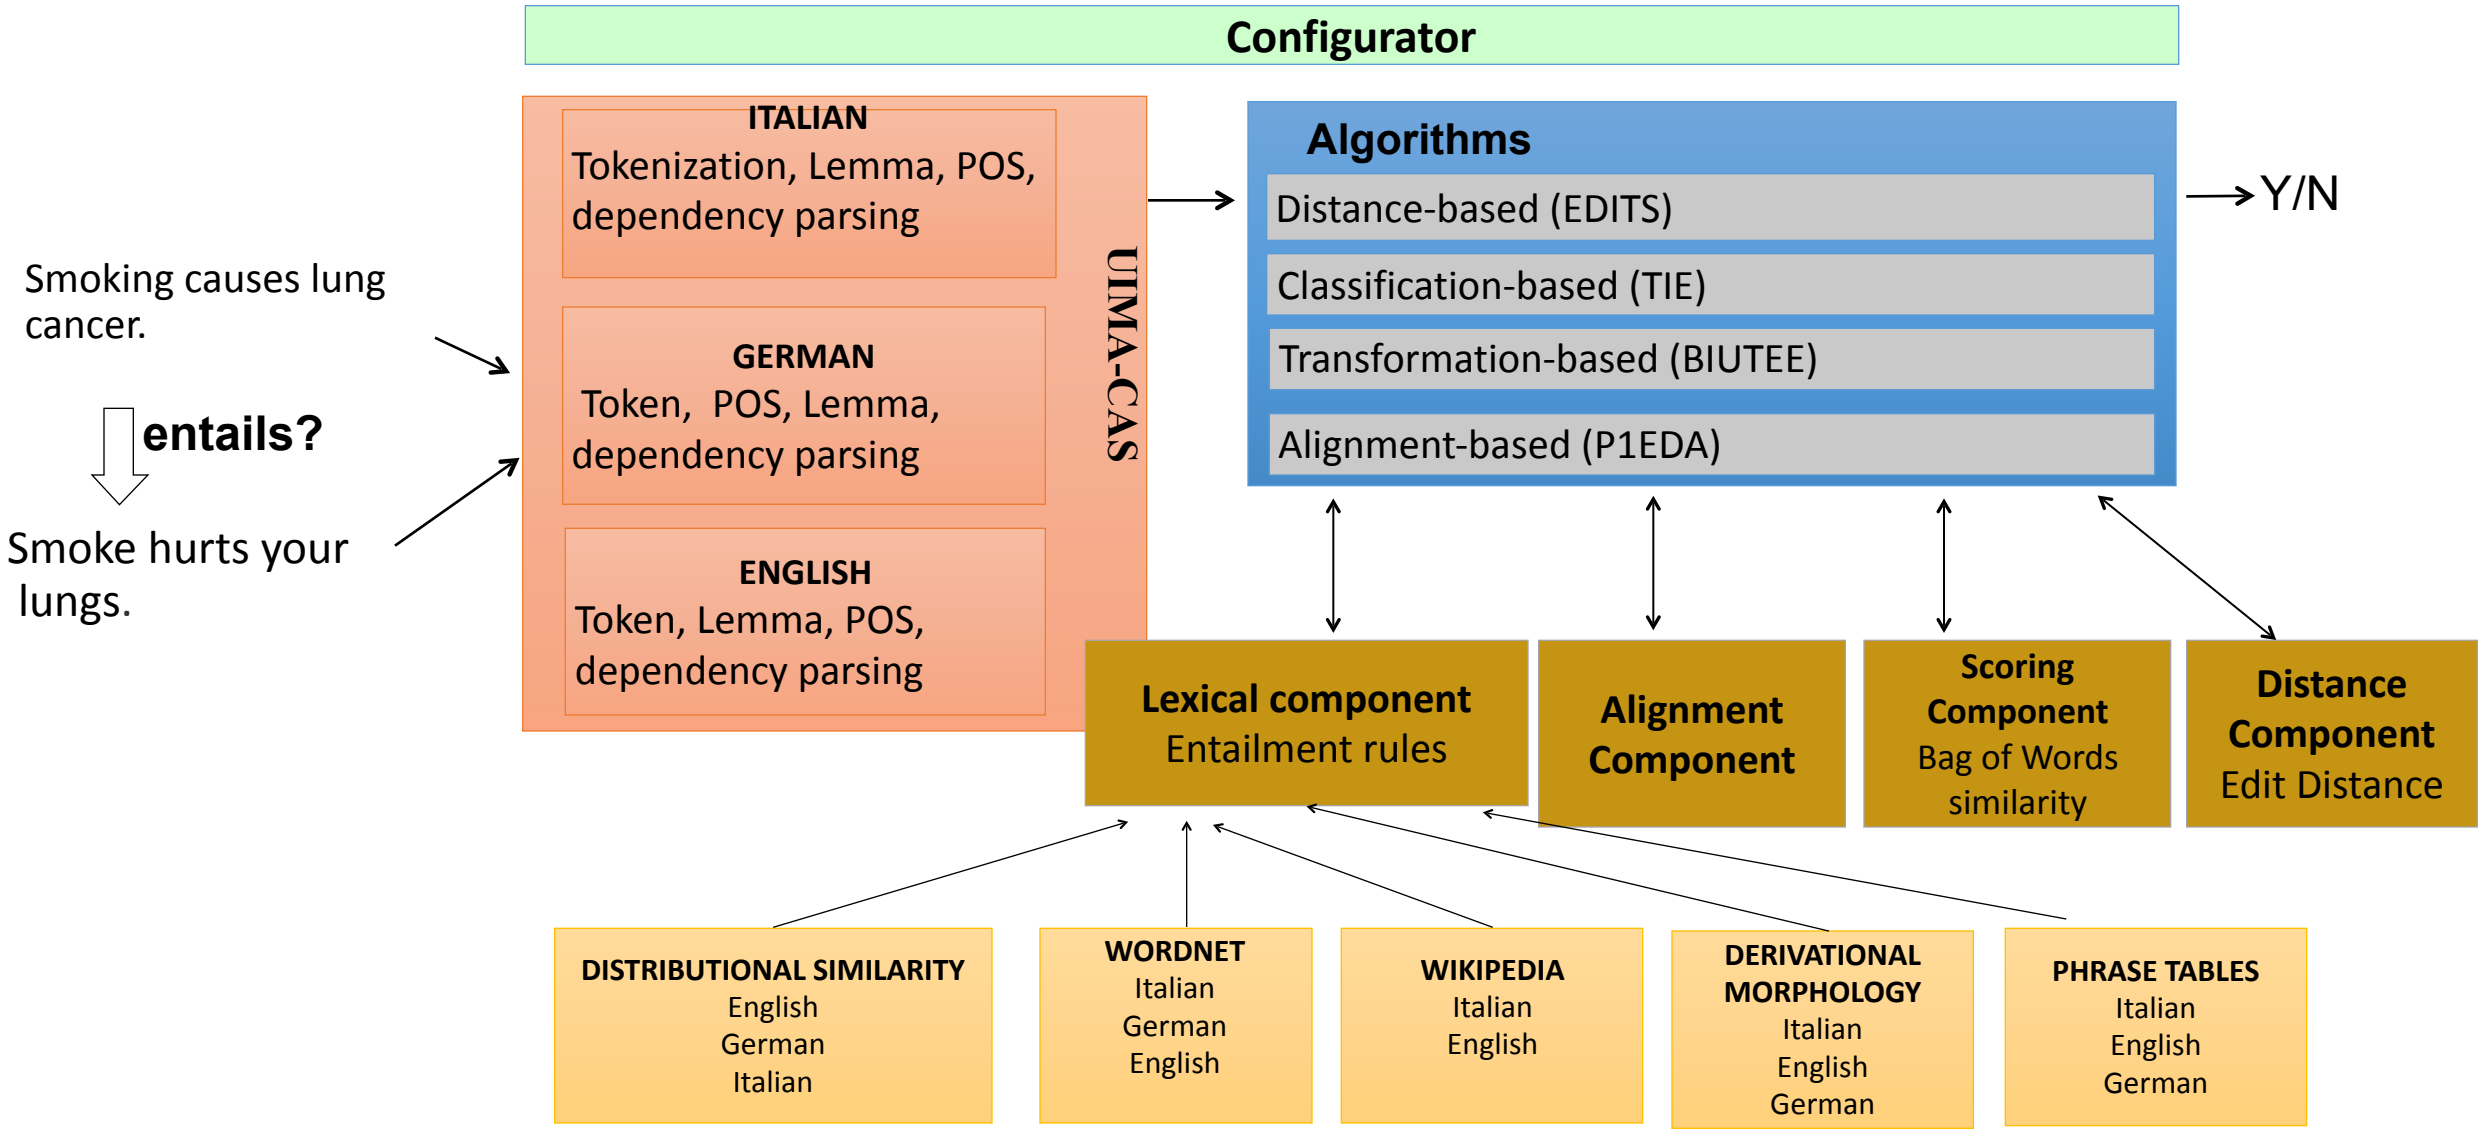

EOP Users

- **Textual Entailment Researchers**
  - Evaluate algorithms to find out their strengths and weaknesses
  - Implement algorithmic ideas
  - Remove influence of resources, preprocessing, ...
  - Extend existing system OR build new system from scratch

- **Textual Entailment End Users**
  - Compare various TE algorithms for applications
  - Does not want to touch code
  - Clear interface (package):
  - Flexible, usable & configurable system
  - Fast prototype to setup simple TE system (Bulgarian)

EOP Distribution

http://hltfbk.github.io/Excitement-Open-Platform/

EOP

Excitement Open Platform for Recognizing Textual Entailment

View project on GitHub

Introduction

The Excitement Open Platform (EOP) is a generic architecture and a comprehensive implementation for textual inference in multiple languages. The platform includes state-of-art algorithms, a large number of knowledge resources, and facilities for experimenting and testing innovative approaches. The EOP is one of the main outcomes of the project EXCITEMENT - EXploring Customer Interactions through Textual EntailMENT. The Platform includes readily-available Recognizing Textual Entailment (RTE) technology, and a modular architecture for text preprocessing, entailment engines and several knowledge resources. The current EOP version covers three languages (i.e. English, German, Italian) and includes tools for creating new resources in other languages. The EOP has been designed to be used in several research use-cases and its distribution package provides a number of utilities.

Recognizing Textual Entailment

Textual Entailment is a directional relation between text fragments. Given two text fragments, one named Text (T) and the other named Hypothesis (H), the Recognizing Textual Entailment task consists in recognizing whether the Hypothesis can be inferred from the Text. We use a graduated definition of entailment: T entails H (T → H) if, typically, a human reading T would infer that H is most likely true. The following is an example of positive entailment:

- Text: If you help the needy, God will reward you.
- Hypothesis: Giving money to a poor man has good consequences.

More on Recognizing Textual Entailment can be found at the ACLwiki Textual Entailment Portal.

Architecture

The EOP takes T-H pairs as input and the output is an entailment judgement, represented by "Entailment" if T entails H, or "NonEntailment" if the relation does not hold. The EOP architecture is based on the concept of modularization with pluggable and replaceable components to enable extension and customization.

The Linguistic Analysis Pipeline (LAP) is a collection of annotation components for Natural Language Processing (NLP) where component integration is based on the Apache UIMA framework. It enables interoperability among components while ensuring language independence.

EOP Demo

Latest Release: v1.1.4

News and Events

October 2014: EOP v1.1.4

May 28-30, 2014: EOP at HLT Village at LREC 2014

Previous Releases

June 2014: EOP v1.1.3
March 2014: EOP v1.1.2
February 2014: EOP v1.1.1

This page was generated by GitHub Pages using the Architect theme by Jason Long.

http://hltfbk.gith

https://github.com/hltfbk/EOP-1.2.0/wiki

EOP RTE

SPIEGEL ONLINE | H2020 | MDFrameWork | JTIG | DFKI | Mathe4ML | Google | IBM | Yahoo | Web | BigData | dict.cc | BMBF - Bekannt...

This repository Search Explore Gist Blog Help gueneumann

hltfbk / EOP-1.2.0

Watch 2 Star 0 Fork 0

Home

Edit New Page

rzanoli edited this page 15 days ago · 8 revisions

The EXCITEMENT Open Platform (EOP) is a generic architecture and a comprehensive implementation for textual inference in multiple languages. The platform includes state-of-art algorithms, a large number of knowledge resources, and facilities for experimenting and testing innovative approaches.

Major changes of release 1.2.0 compared to the previous release 1.1.4

New Features:

- P1EDA (alignment-based EDA) – this EDA tries to align T with H, according to different aspects of their representation: lexical, syntactic, etc. It eventually classifies whether or not T entails H by considering the plausibility of those alignments.
- LAP for Bulgarian language.

Bug Fixes:

- Wrong Ant build script and documentation for installing TreeTagger.

Pages 4

Documentation

- Licence
- Requirements
- Installation
- Quick Start
- Step by Step Tutorial
- Entailment Algorithms
  - BIUTEE
  - EDITS
  - TIE
  - P1EDA
- Lexical Resources
- Configuration Files
- FAQ

Get Involved

- How
- Mailing Lists
- Available Projects
- Results Archive

Clone this wiki locally

https://github.com/hltfbk/EOP

Clone in Desktop

Open Source Distribution of EOP

- Quick Code Integration
  - Git, Github, Maven, Jenkins
- Quality Control
  - Code quality tools (e.g. check style, find bugs)
- Additional Highlights
  - Archive for Experiments
  - GitHub wiki pages (release-specific documentation)
  - Two Distributions: API and Command Line Interface
- License: General Public License (GPL) version 3

Open Source Distribution of EOP

- Quick Code Integration
  - Git, Github, Maven, Jenkins
- Quality Control
  - Code quality tools (e.g. check style, find bugs)
- Additional Highlights
  - Archive for Experiments
  - GitHub wiki pages (release-specific documentation)
  - Two Distributions: API and Command Line Interface
- License: General Public License (GPL) version 3

EXPERIMENT: T5

Configuration ID: MaxEntClassificationEDA_Base+WN+TP+TPPos+TS_EN

Data Set: RTE-3
Language: EN
Lexical Resources: WN,TP,TPPos,TS
Preprocessing: MaltParserEN
Results(Accuracy): 0.65250

Author name: Günter Neumann
Affiliation: Deutsche Forschungszentrum für Künstliche Intelligenz GmbH (DFKI)
Contact Information: neumann [@] dfki.de

Download the configuration file, model and results

Overview – Release Management

- Keeping several code versions (master branch, releases)
- Automatic methods for
  - creating new releases and resource distributions
  - maintenance of release-specific documentation
  - Generating Web Page (EOP web site)
- Separate documentations for end users and developers

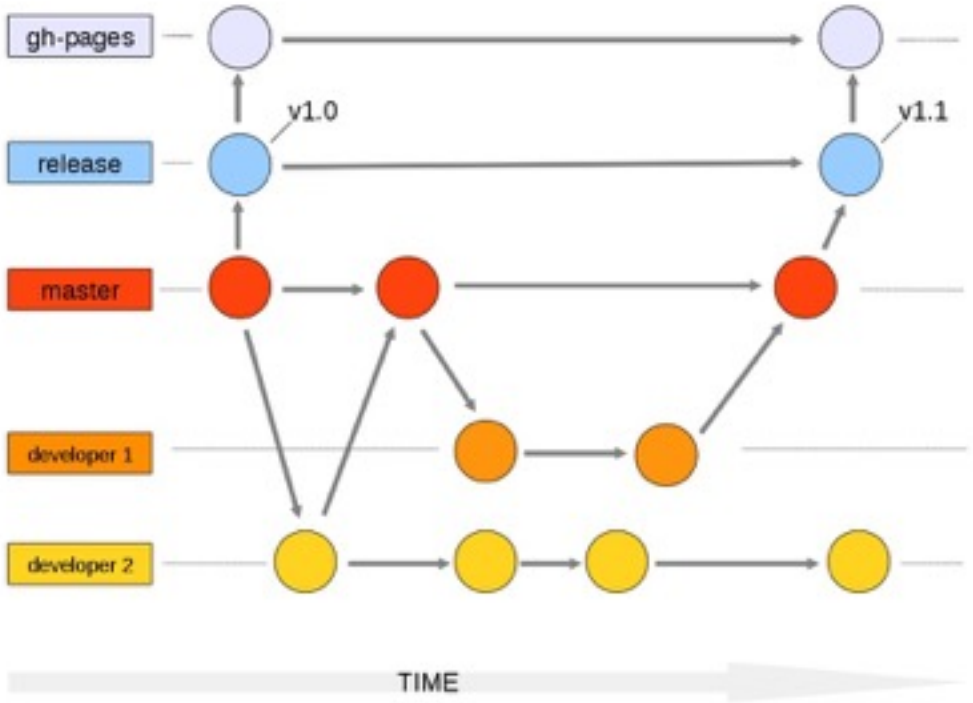

EOP in Numbers (08/09/2014)

- EOP GitHub repository:
  - 52 Members (people who forked the EOP Repository)
- Mailing lists:
  - developers: 21
  - users: 24 (12 external users)
- EOP v1.1.3
  - Downloads: 77
  - Experiments Archive: 13 experiments
    - 96 experiments in the current developers version EOP v1.1.5
  - Download + Installation: 10 min by a shell script

Learn More

- EXCITEMENT project web site: http://www.excitement-project.org

- B. Magnini, R. Zanoli, I. Dagan, K. Eichler, G. Neumann, T.-Gil. Noh, S. Pado, A. Stern, O. Levy: The Excitement Open Platform for Textual Inferences. In proceedings of ACL demo session, June 2014.

- S. Pado, T.-G. Noh, A. Stern, R. Wang, R. Zanoli: Design and Realization of a Modular Architecture for Textual Entailment. Natural Language Engineering. Cambridge University Press, 2014.

- T.-G. Noh, S. Pado. Using UIMA to structure an Open Platform for Textual Entailment. 2013. Proceedings of the UIMA@GSCL workshop.

Building Entailment Graphs

Customer Interactions Scenario

Int-448:

Efficient service. Quick through security and check in. Staff could have been a bit more friendly though and leg room in standard class was quite poor.

Int-202:

Everything ran smoothly and well. Only complaint is lack of leg room with seating with tables. Very cramped when all seats are taken.

Int-275:

The leg room in economy class is not enough I was constantly being kicked by opposite passenger I travel by train lots and this compares badly to other trains

Int-303:

My only gripes, not enough leg room in standard and I think it would be chic to have refreshments served in carriages , either trolley or trays like in theatres .

EXCITEMENT application scenario

Requirements

- Need for customer interaction analytics
  - **Compact** representation (show just relevant information)
  - **Informative** representation: general categories (e.g. “food”, “internet”) are not enough
- Need to manage streams of data
- Multiple channels: e-mail, speech, social media
- Noisy data: automatic transcriptions, social media style, etc.
- Multiple languages
  - Excitement: English, Italian, German

Challenge

- Core technology: entailment graphs based on the EOP platform
- Current experiments based on the Alignement-based algorithm

Extracting Fragments from Interactions

TOPIC: Reasons for dissatisfaction in railway service

Int-448: Efficient service. Quick through security and check in. But **leg room in standard class was quite poor**.

Int-202: Everything ran smoothly and well. Only complaint is **lack of leg room with seating with tables**.

Int-275: **Seating is very cramped** – my journey has been very uncomfortable with the person next to me taking up most of the space we have.

Int-303: My only gripes r **not enough leg room in standard** and I think it would be chic to have refreshments served in carriages , either trolley or trays like in theatres .

Building Fragment Graphs

Int-448

Building Fragment Graphs

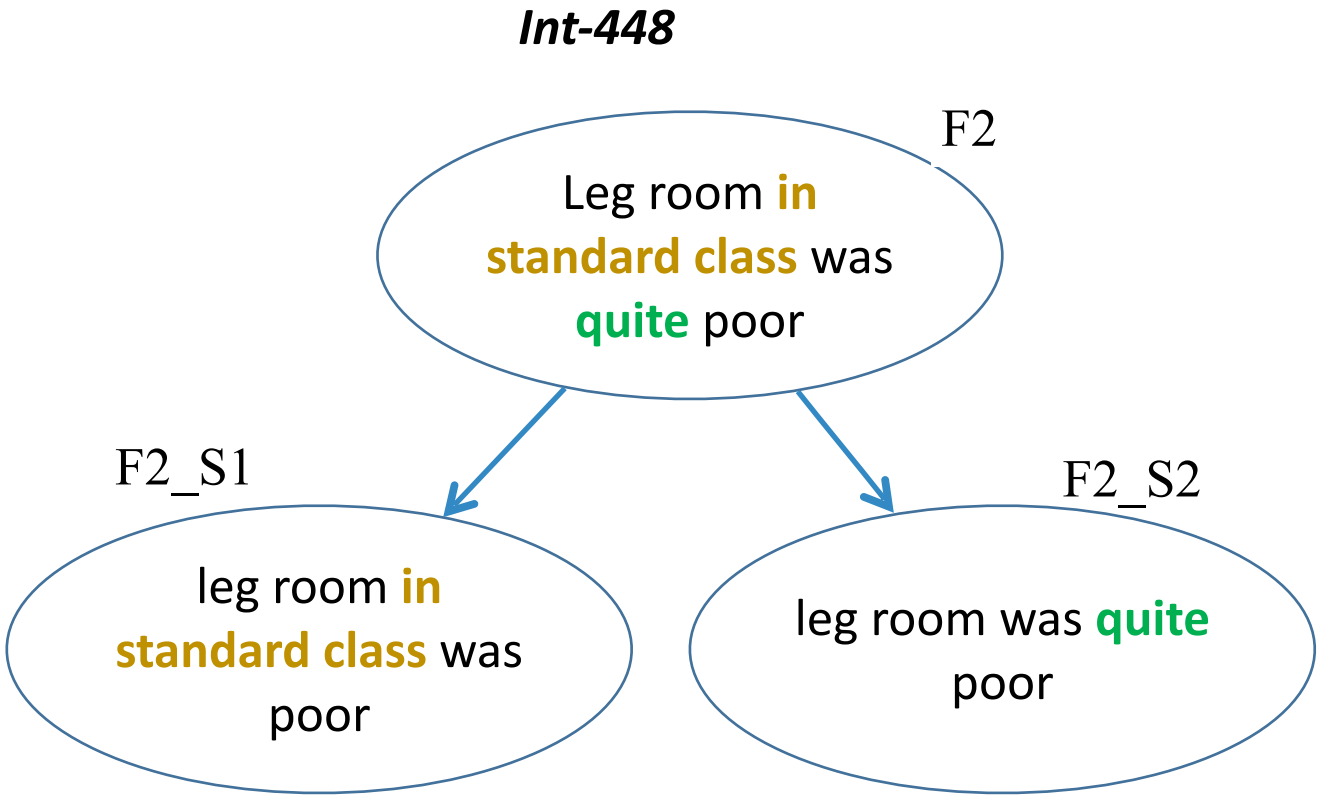

Building Fragment Graphs

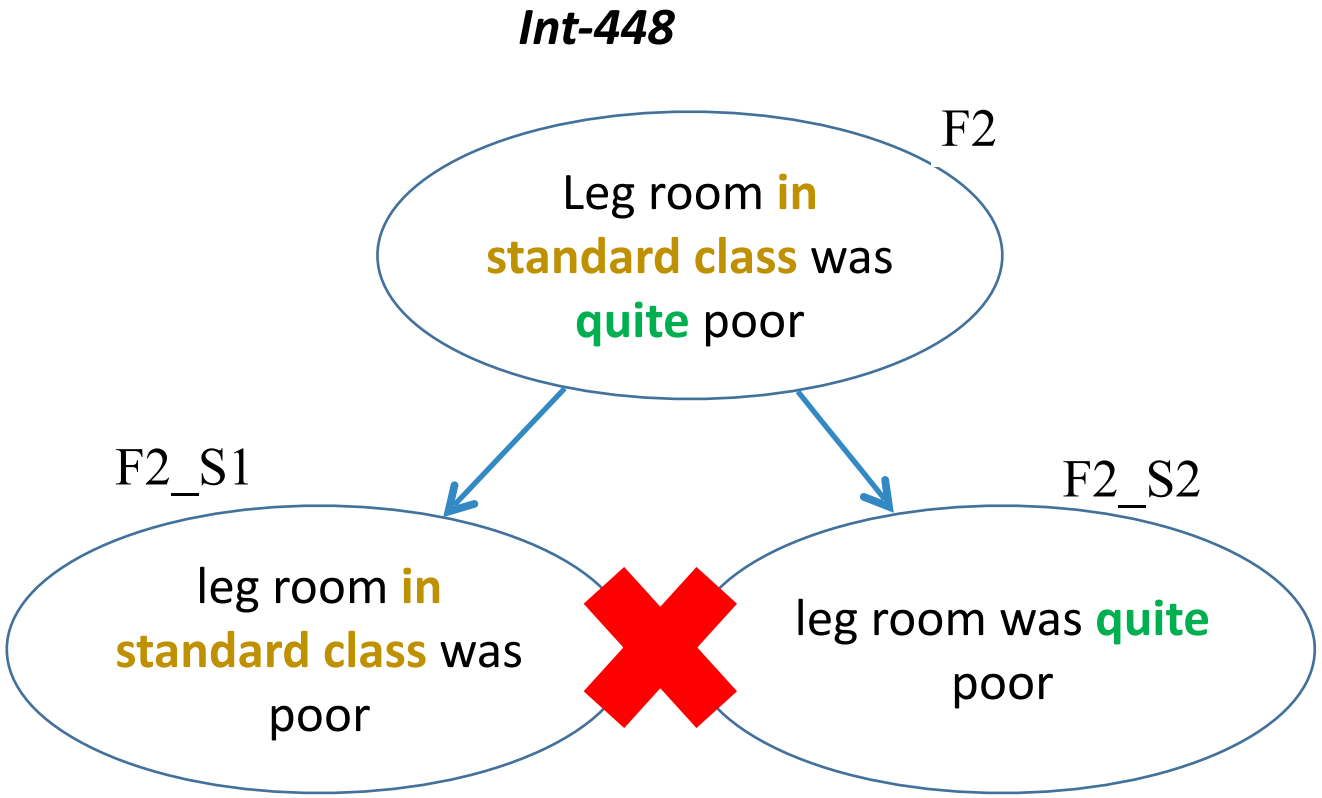

Building Fragment Graphs

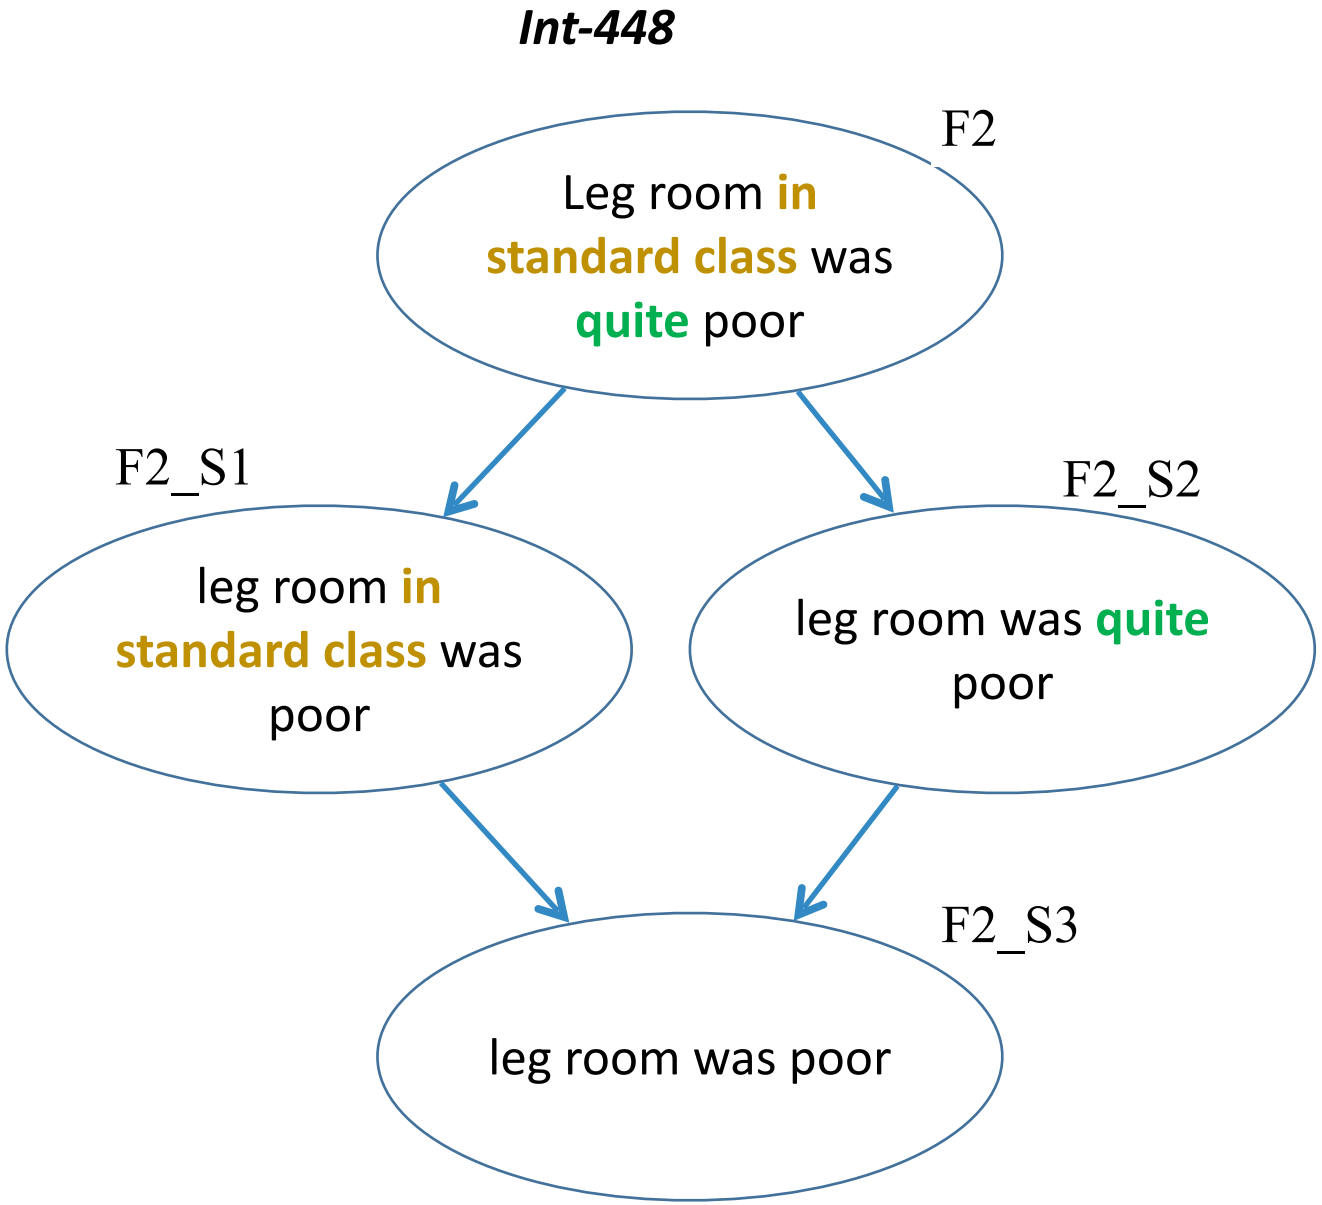

Building Fragment Graphs

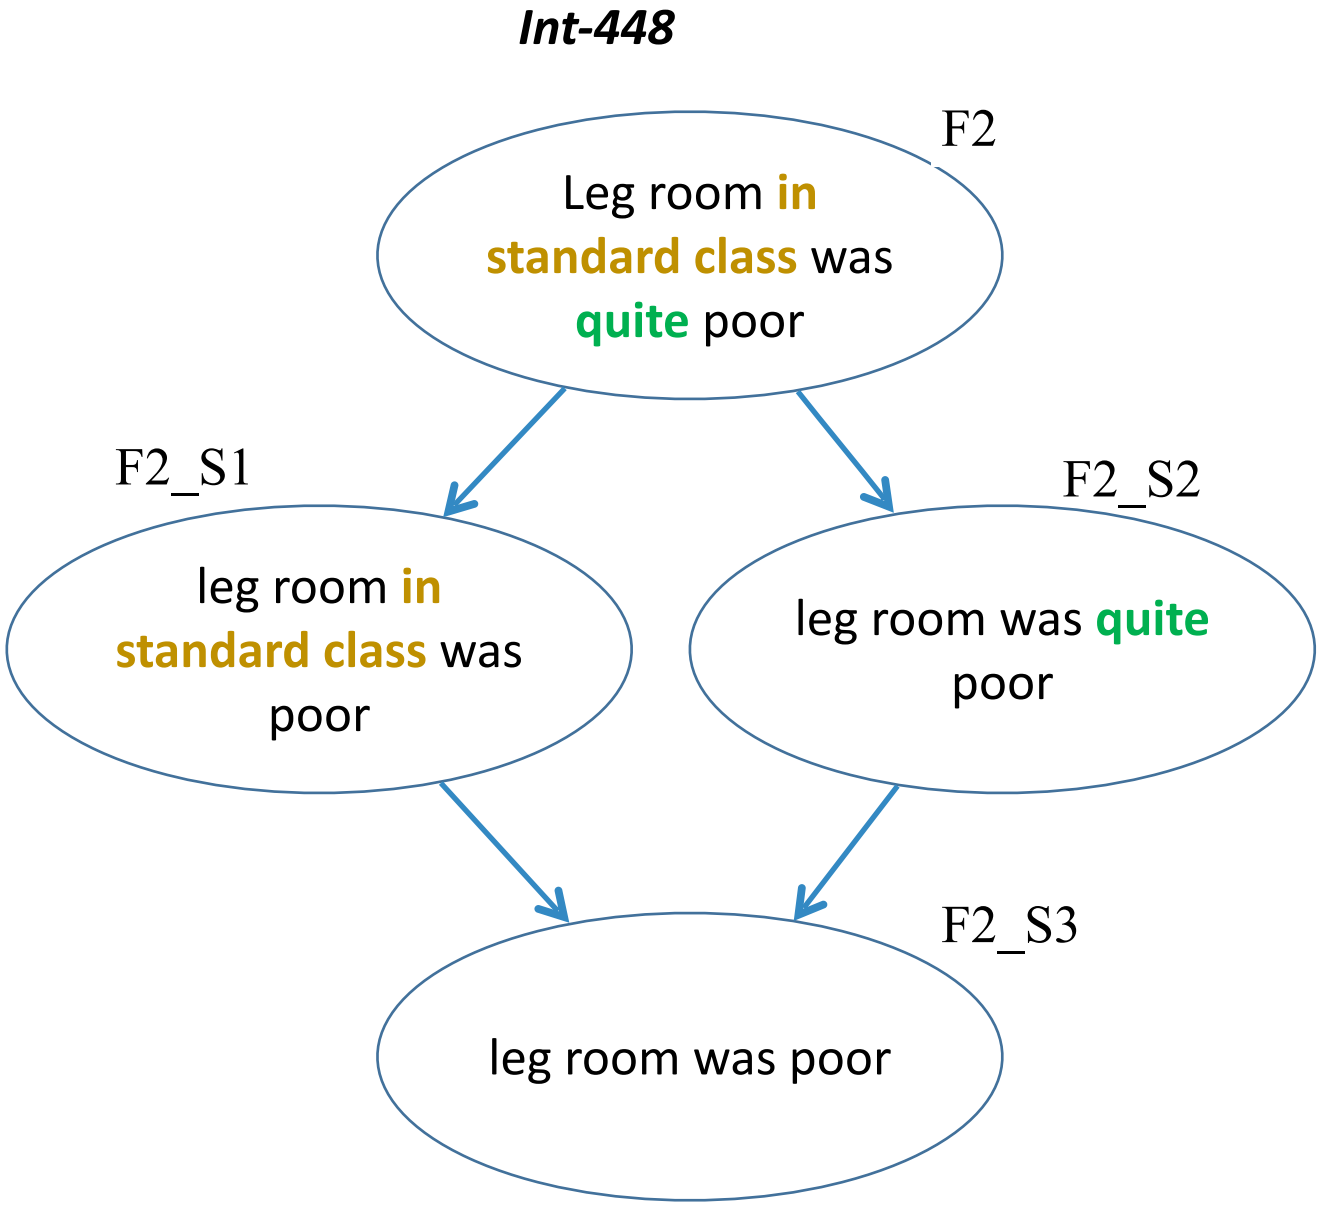

Result: a DAG

- rooted in Fragment
- Base predicate (fragment without all modifiers as only leaf)

Merging Graphs with the EOP

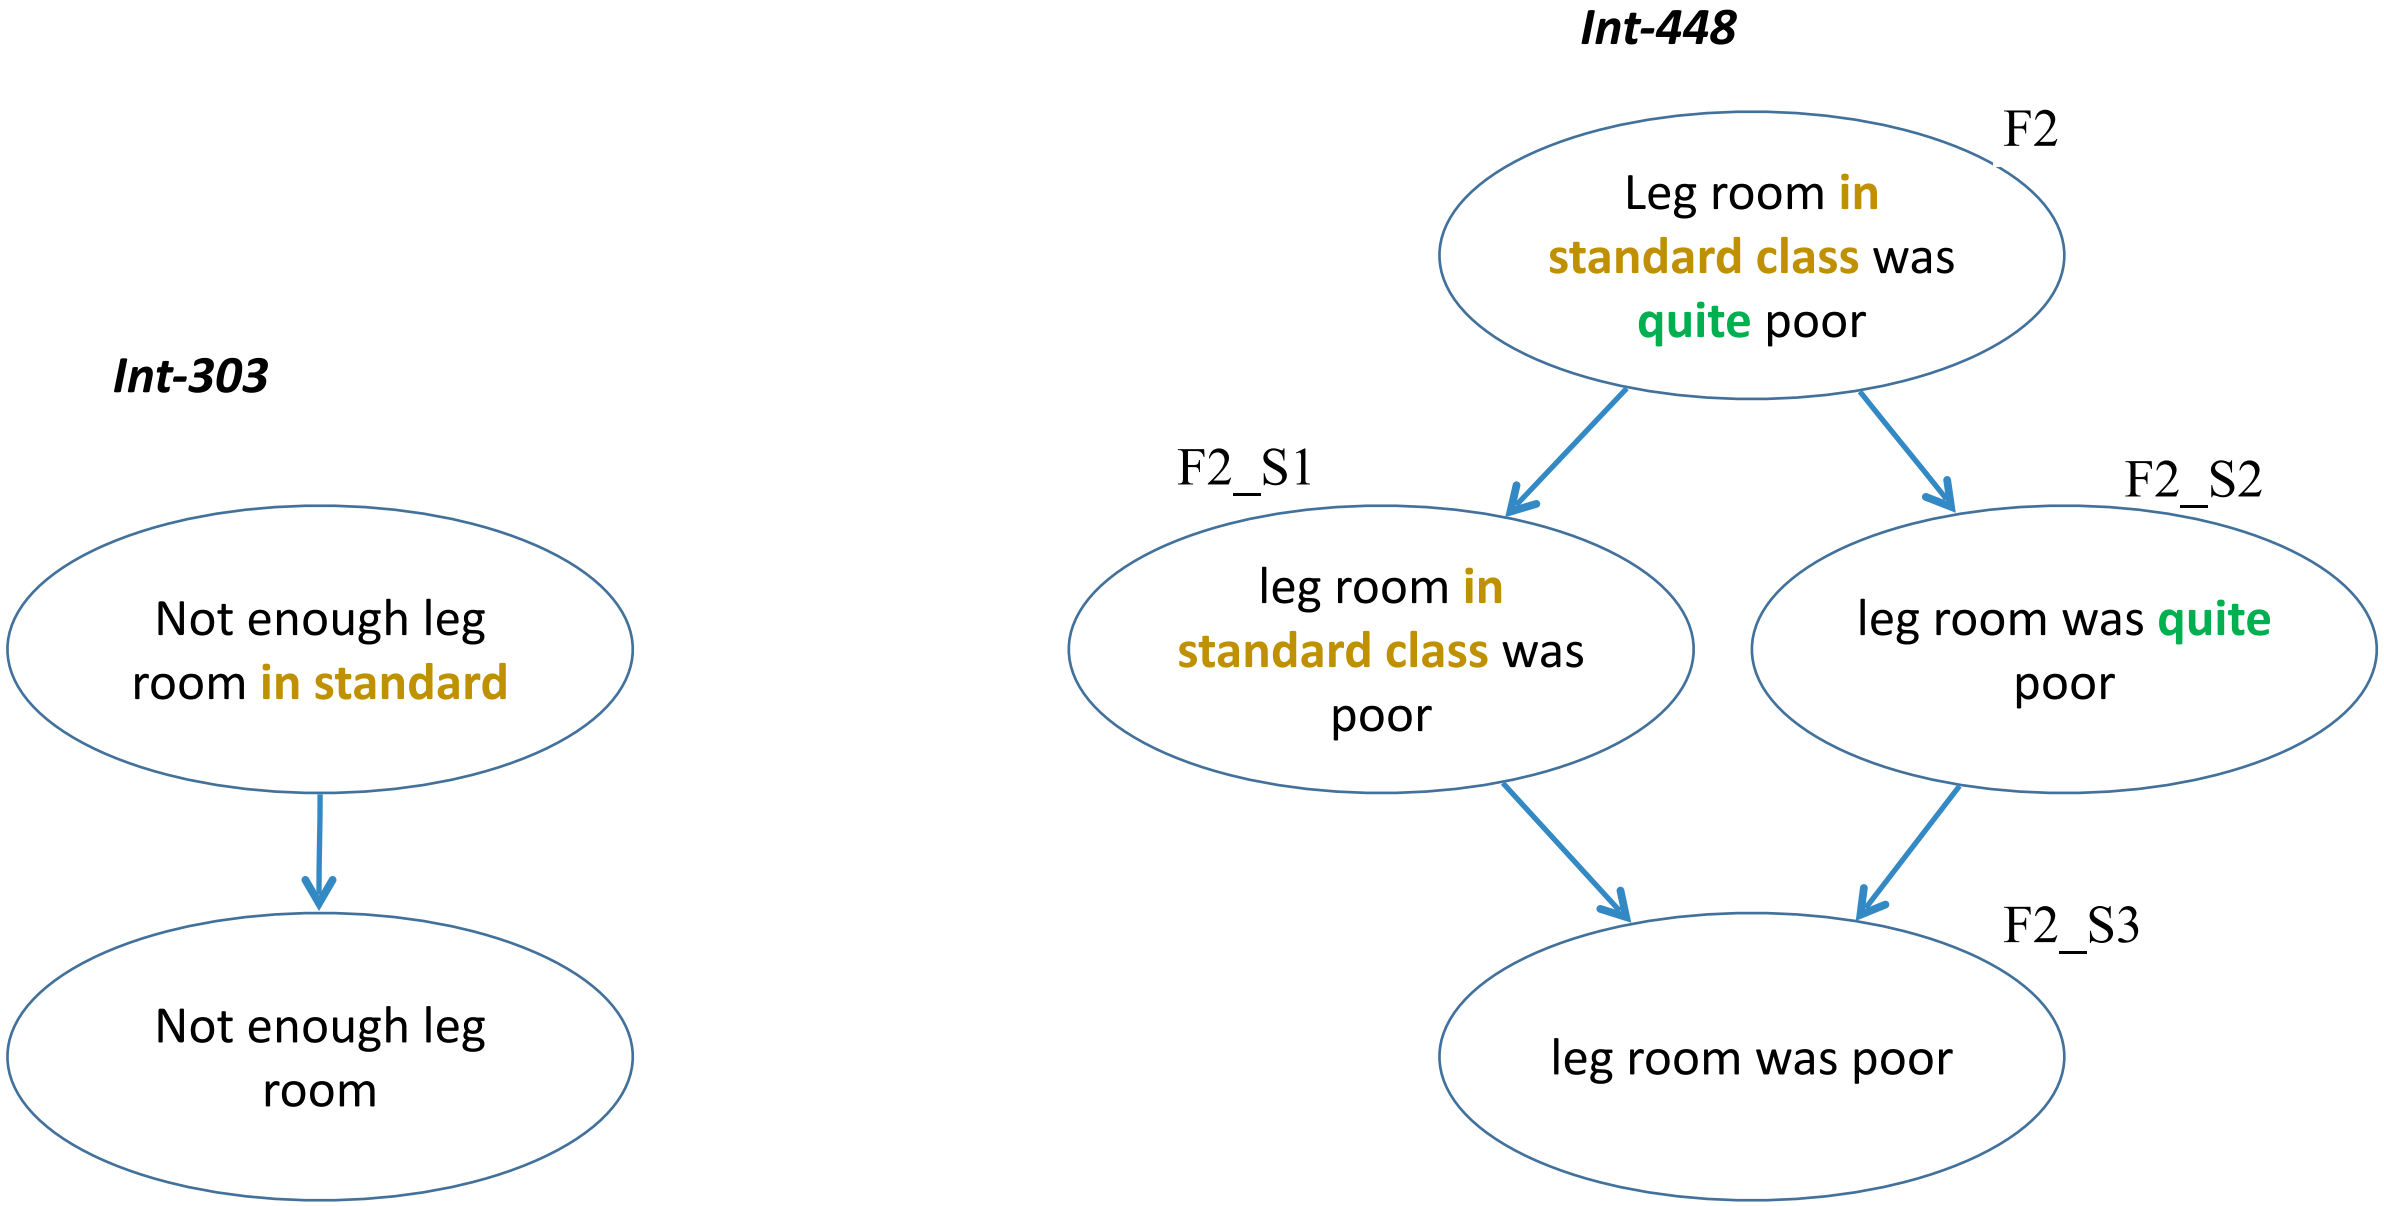

Merging Graphs with the EOP

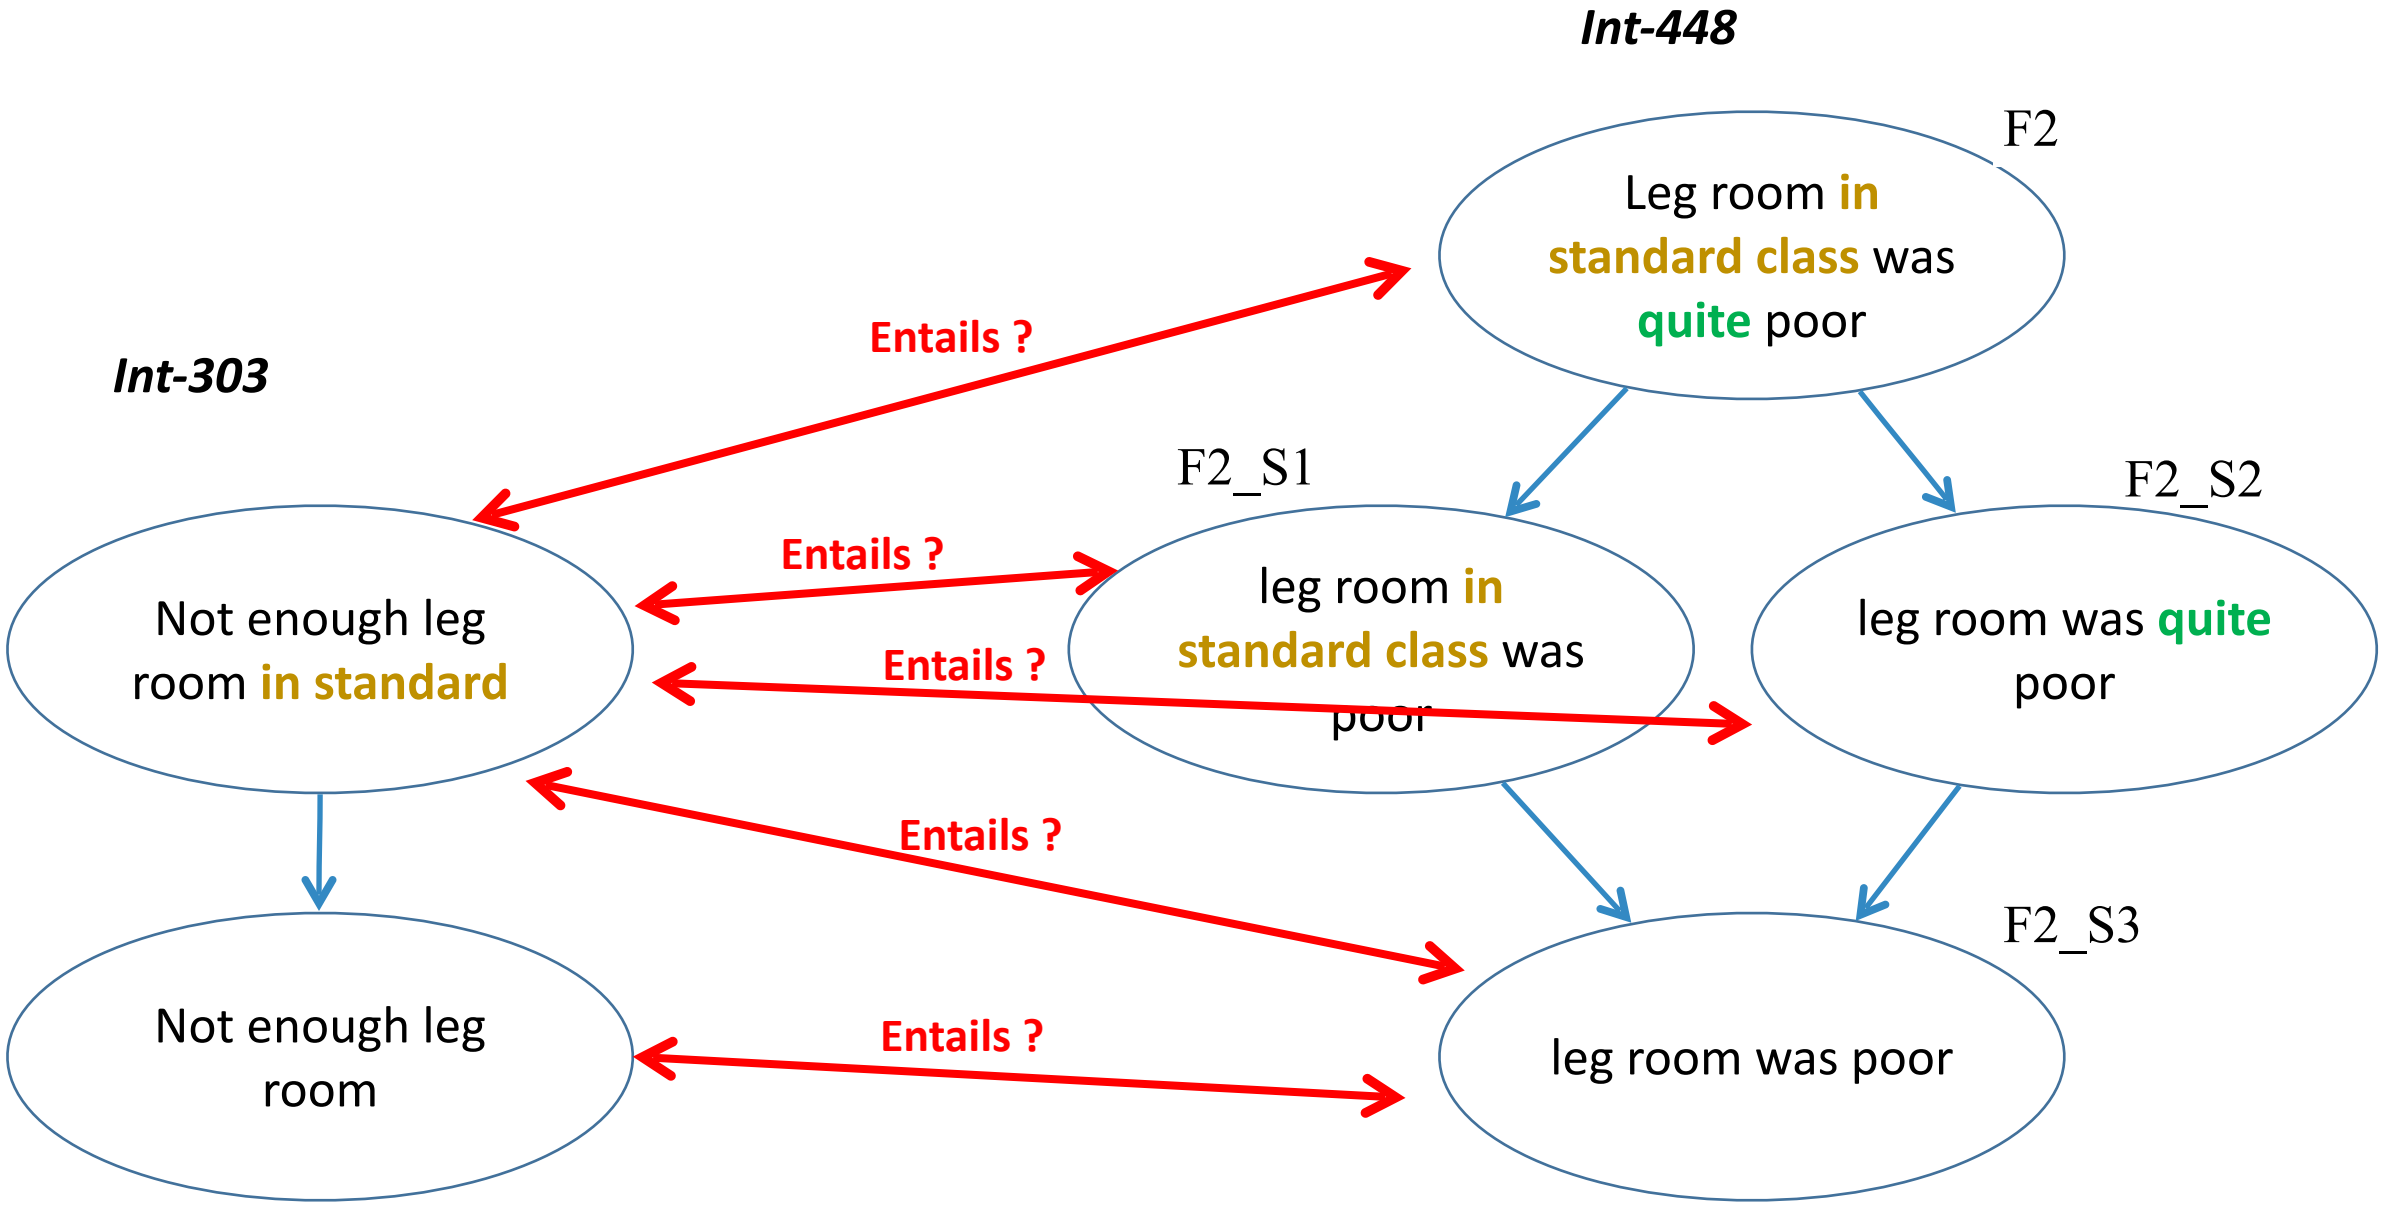

Merging Graphs with the EOP

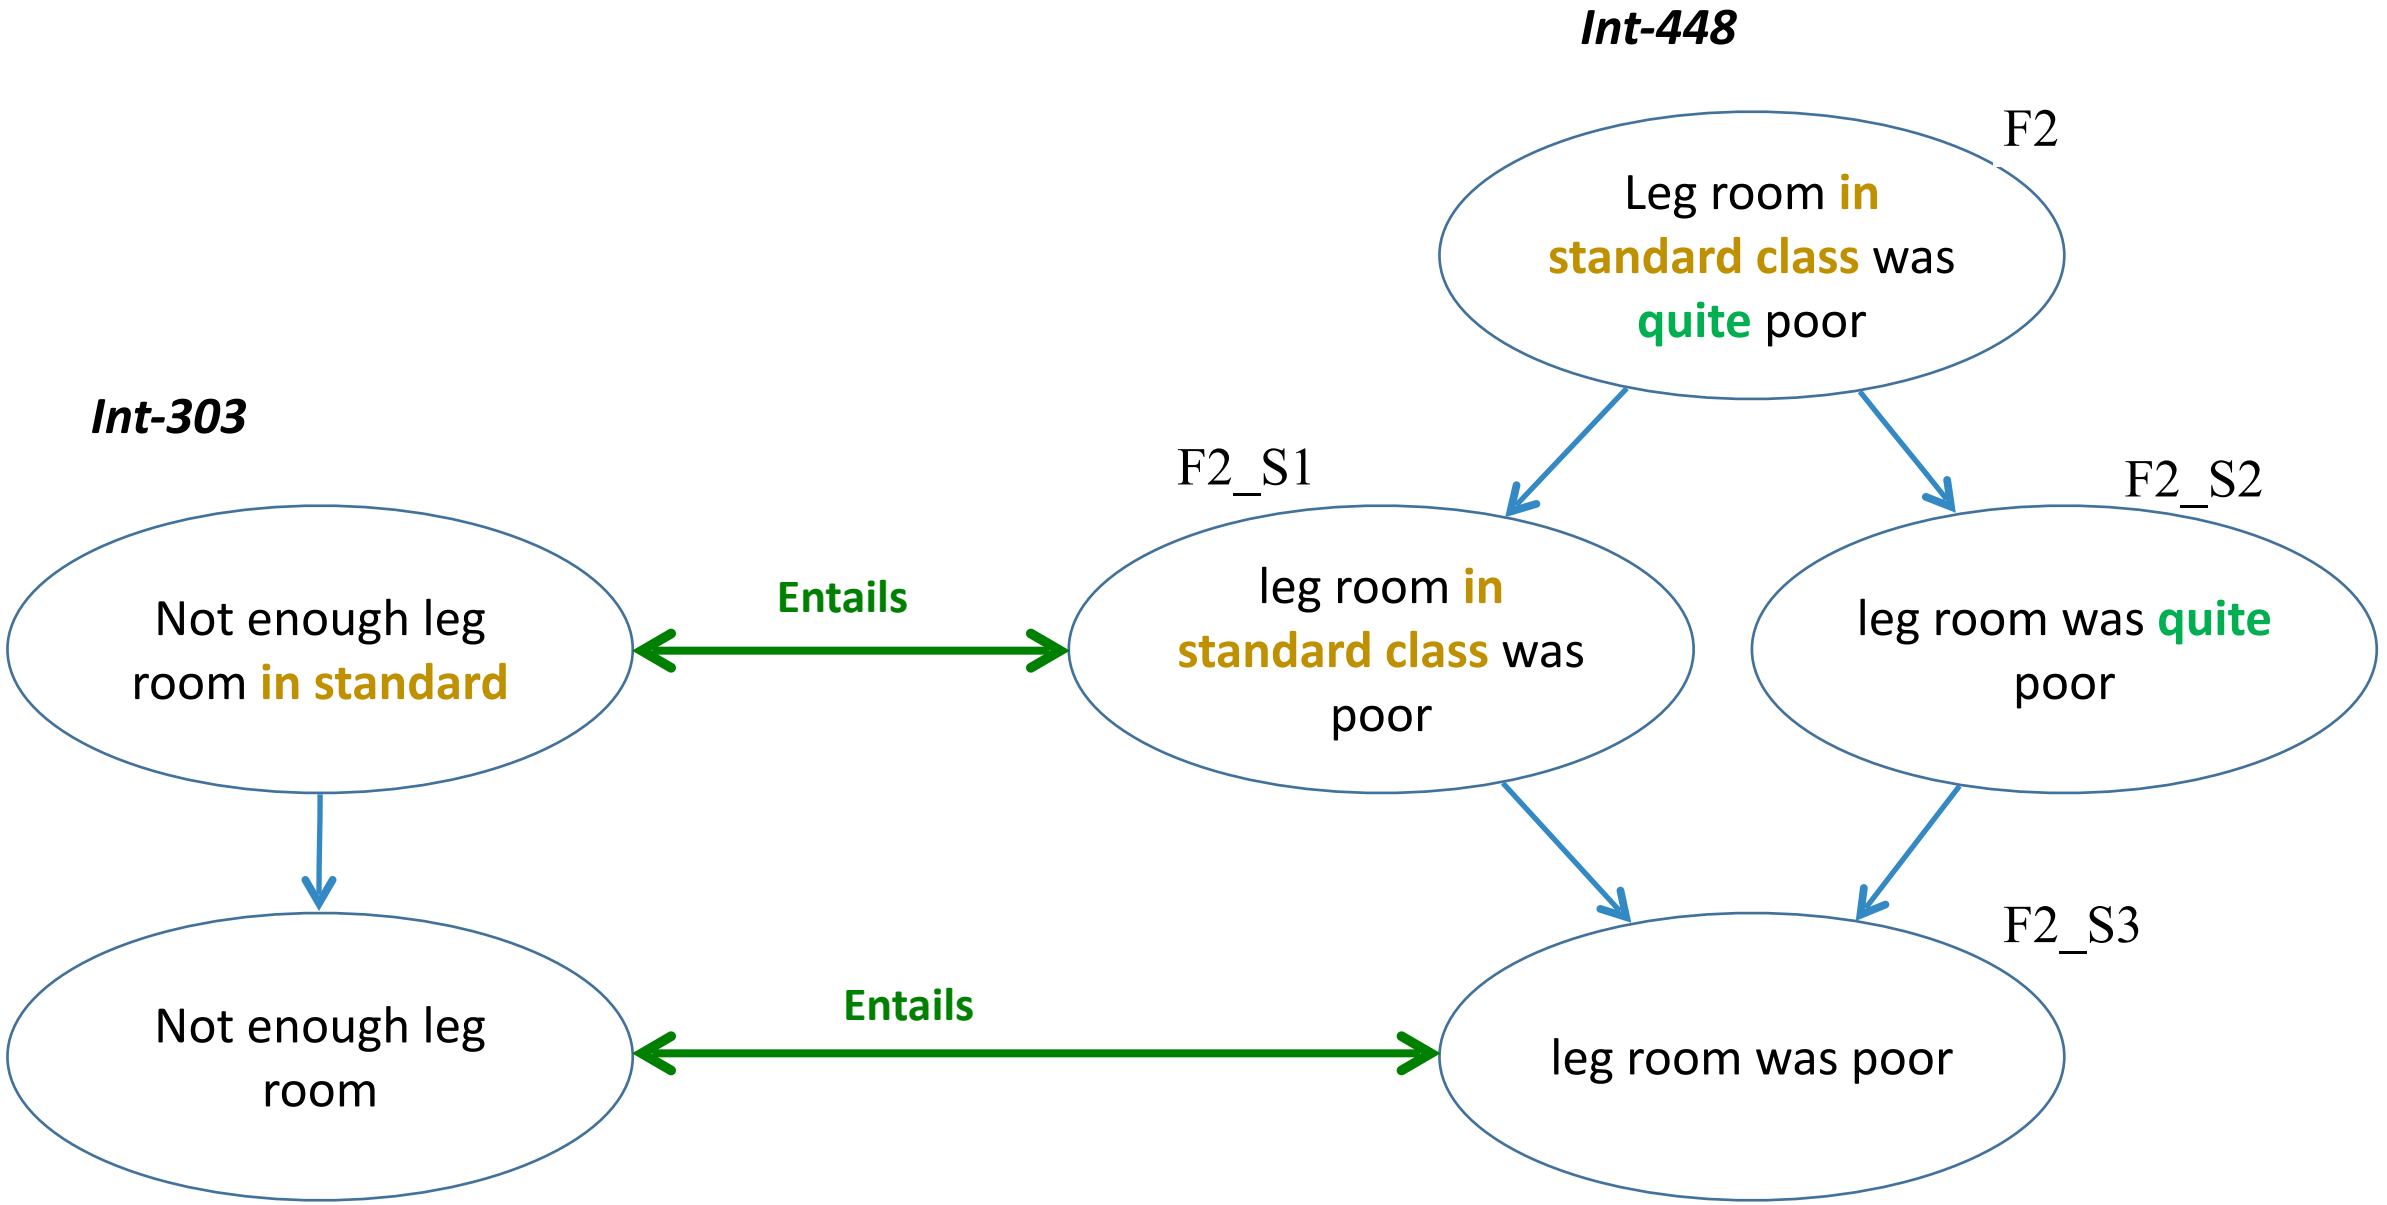

Merging Graphs with the EOP

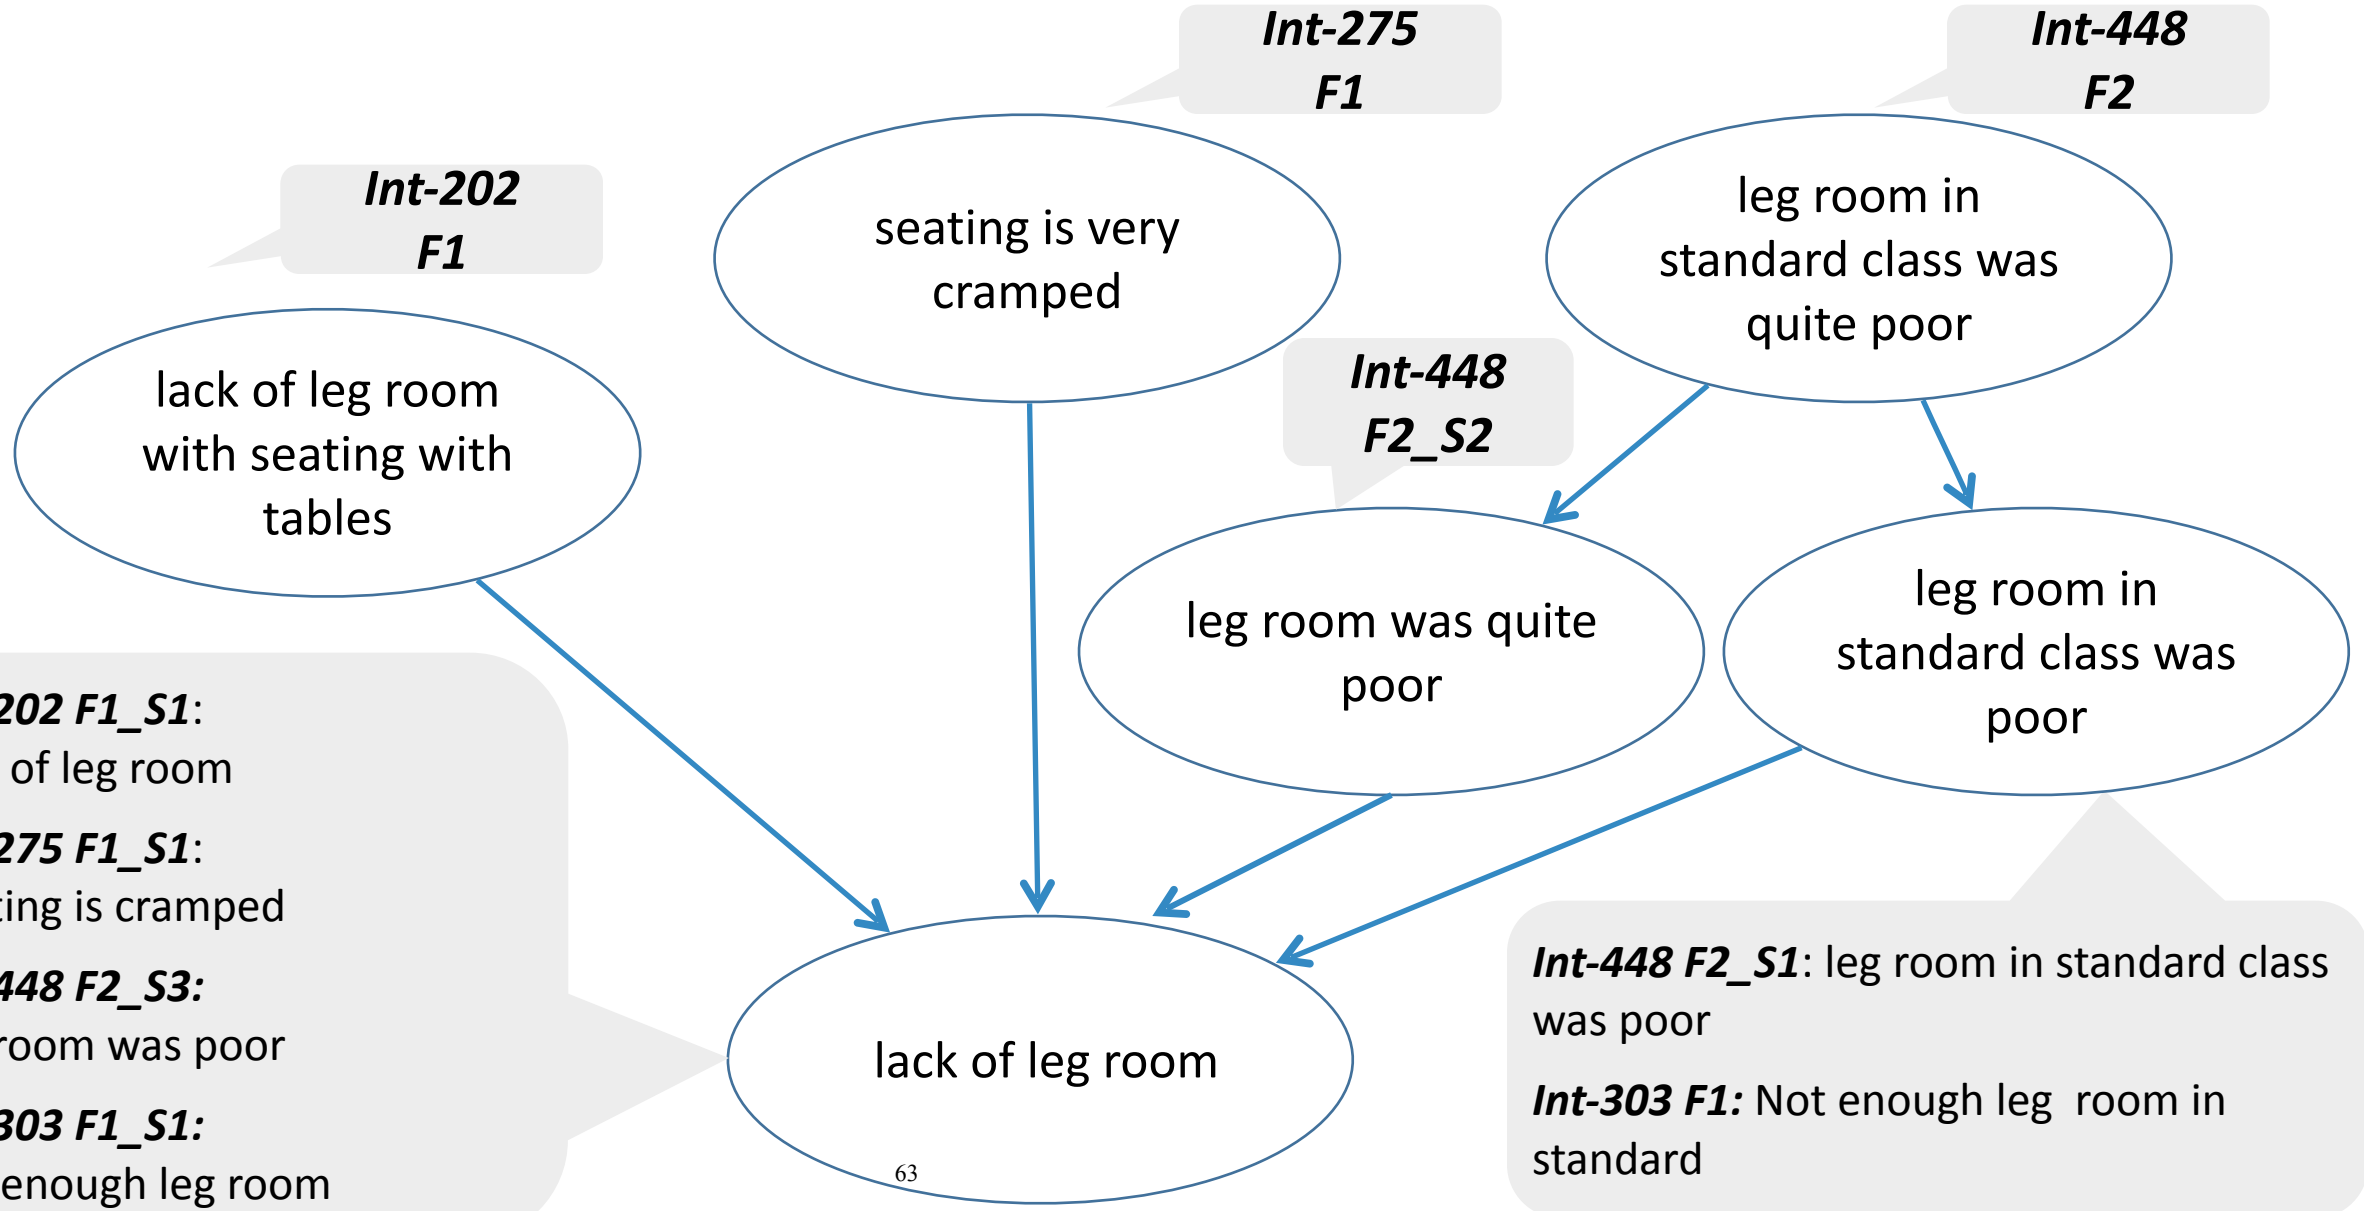

Conclusion

- **Textual Entailment** provides a generic perspective for inference over textual expressions
- **Textual inference technology** is still in early stages, with limited yet potentially useful performance
- The **EXCITEMENT Open Platform** offers available technology for research
- **Entailment Graphs** have a potential for text exploration applications
- **Datasets and baseline results** for customer interactions are available for further research