



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,
Saarbrücken

Sebastian Pado

University of Heidelberg

Excitement project

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

One text **implies the meaning**
of the other

pepper may trigger sneezing



allergies can be produced by hot spices

What is applied *textual inference*?

“Match” different text fragments where:

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

paraphrasing
bi-directional entailment

pepper may trigger sneezing



pepper can cause sneezing

One text **implies the meaning**
of the other

(directional) textual entailment

pepper may trigger sneezing



allergies can be produced by hot spices

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

Search

allergenic foods

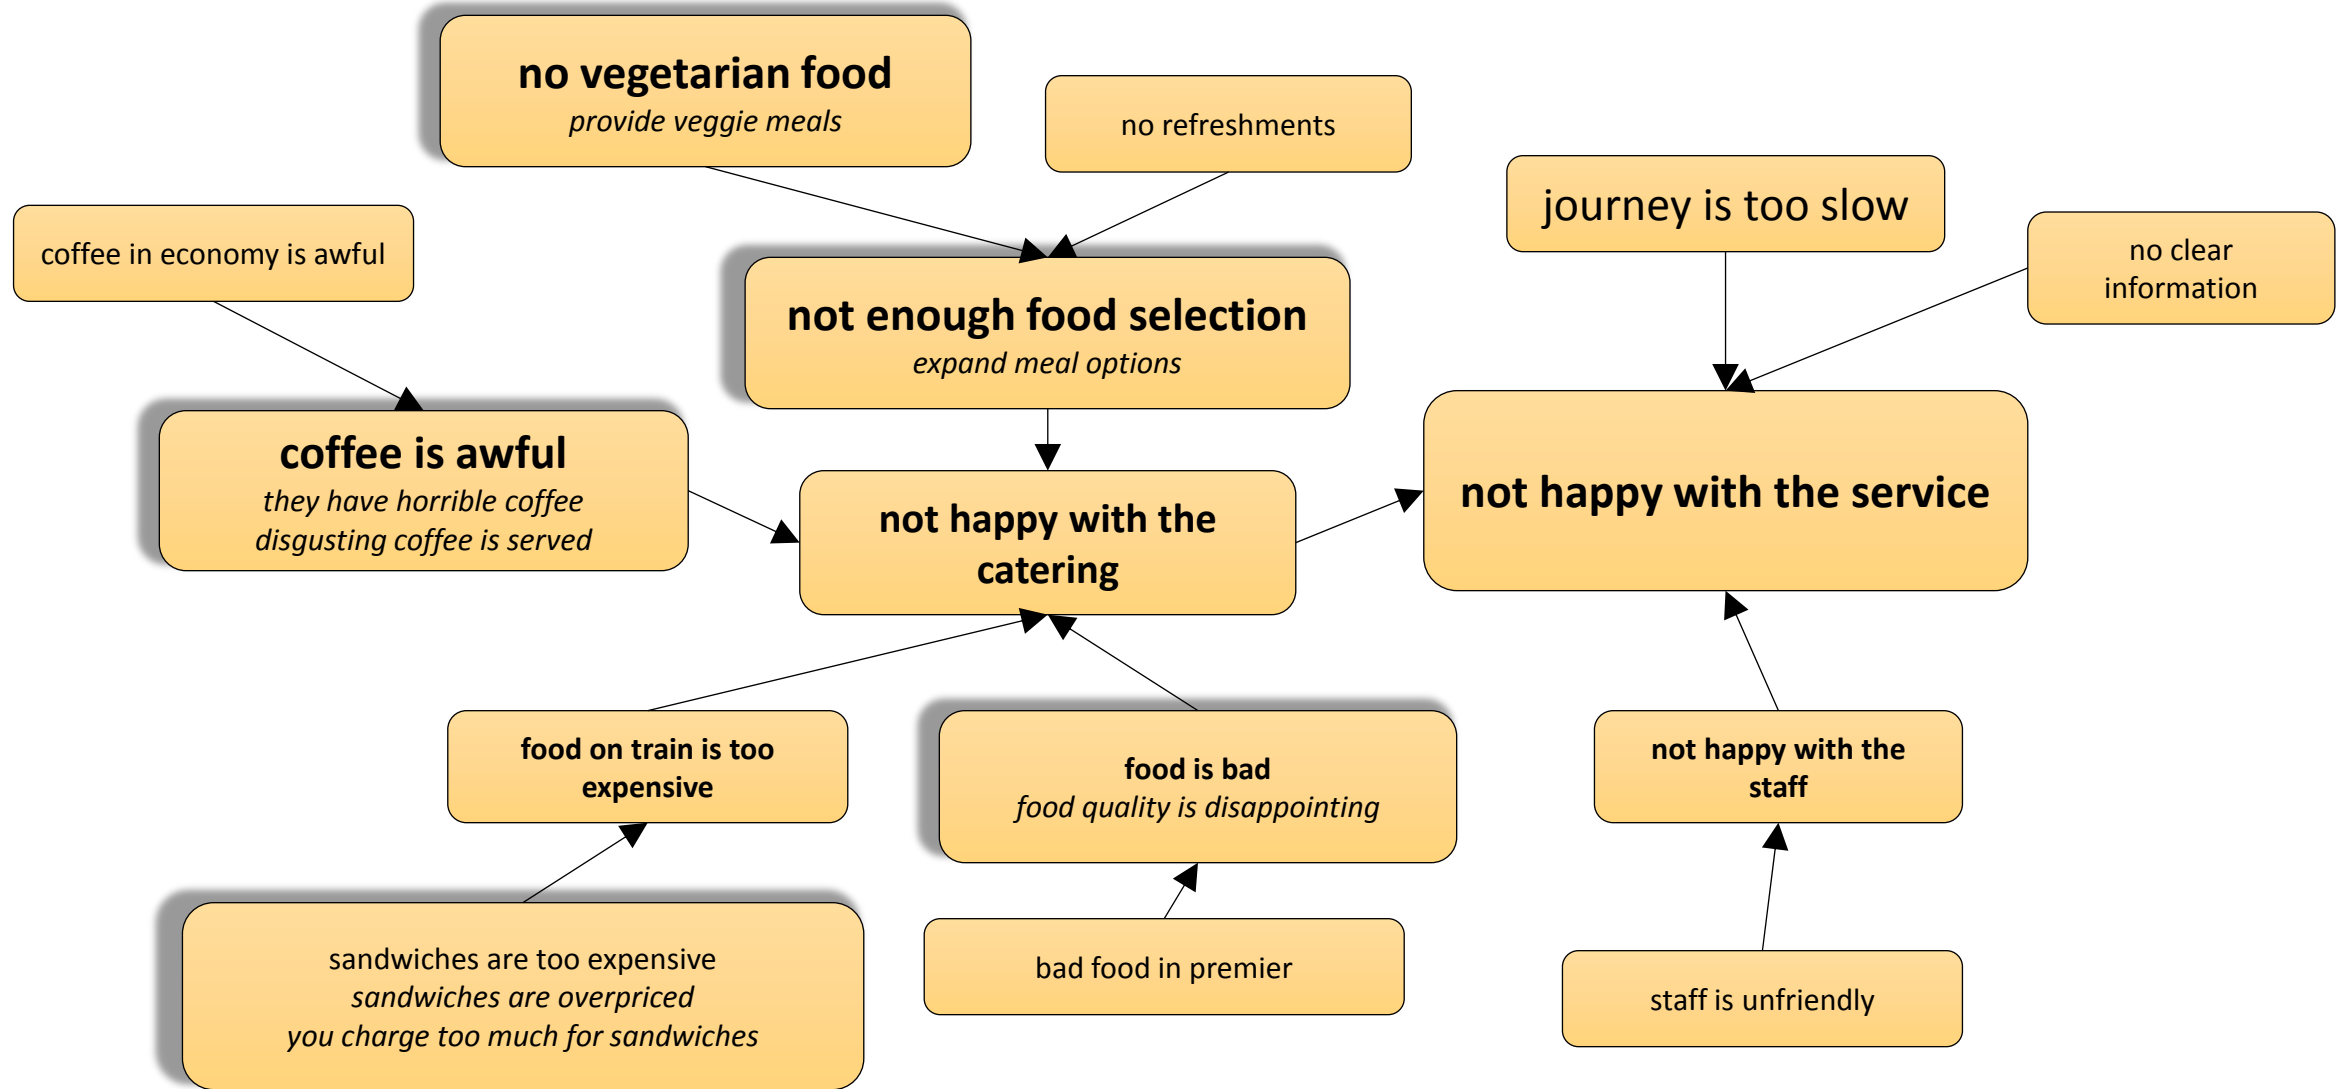
Information Extraction

Extract pairs of foods and symptoms

Summarization

Summarize documents about allergies

Novel Application: Text Exploration



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



P: Peter was Susan's husband

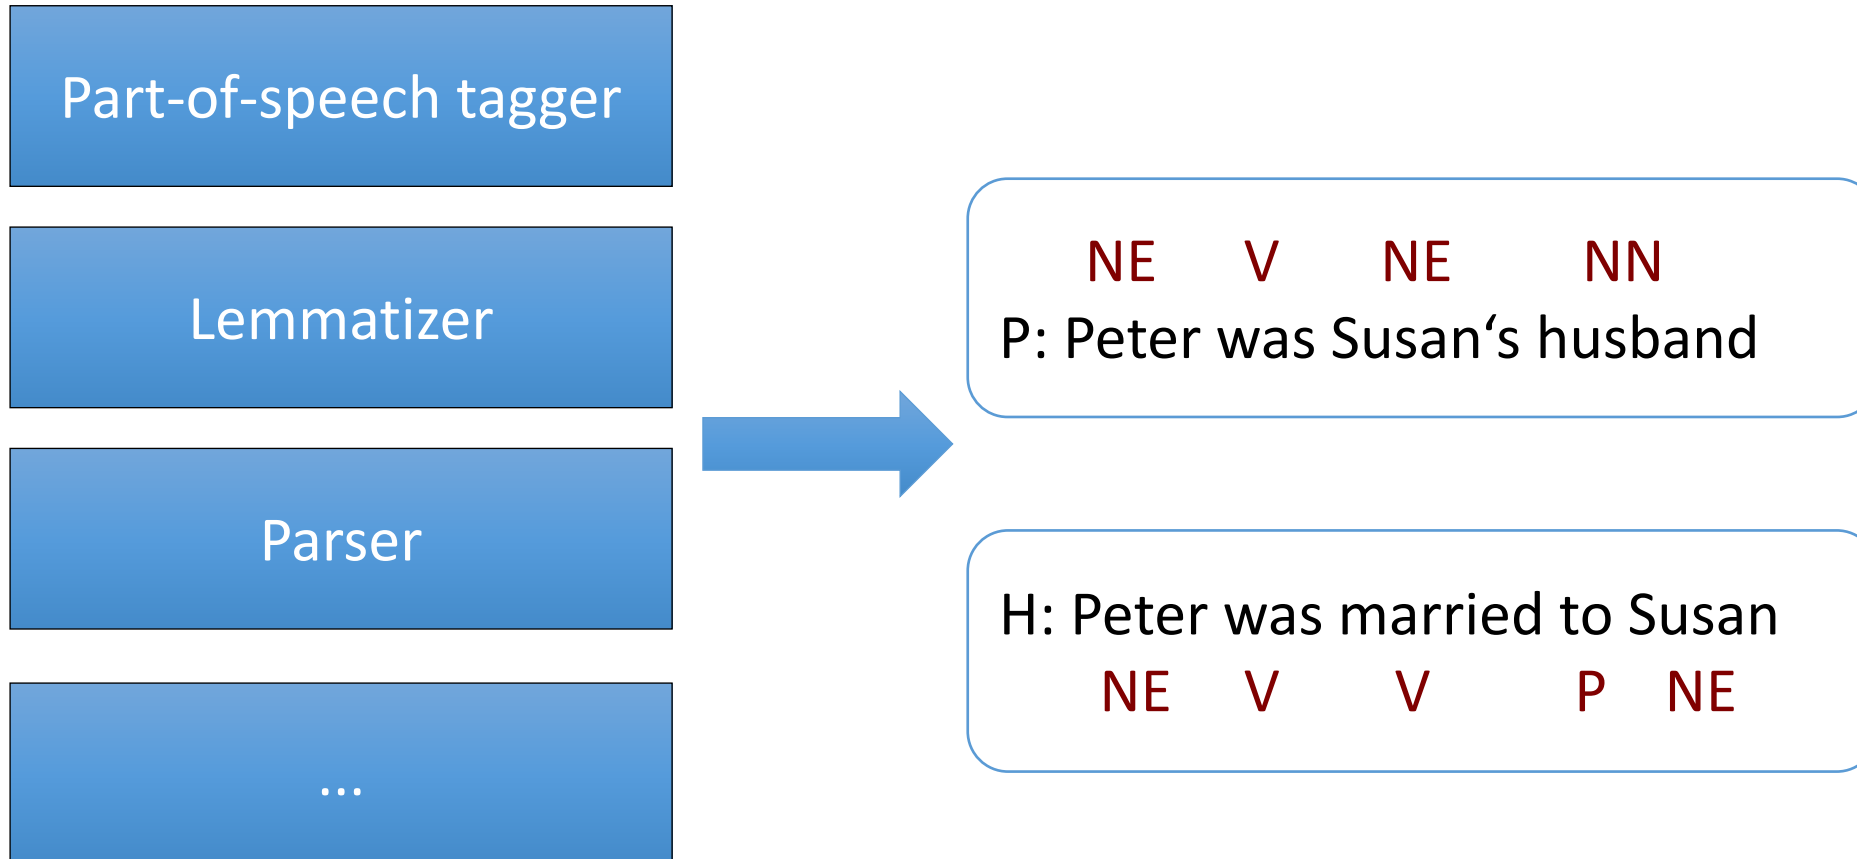
H: Peter was married to Susan

P: Peter did not know Susan

H: Peter was married to Susan


Alignment-based Entailment: The Algorithmic Level

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



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

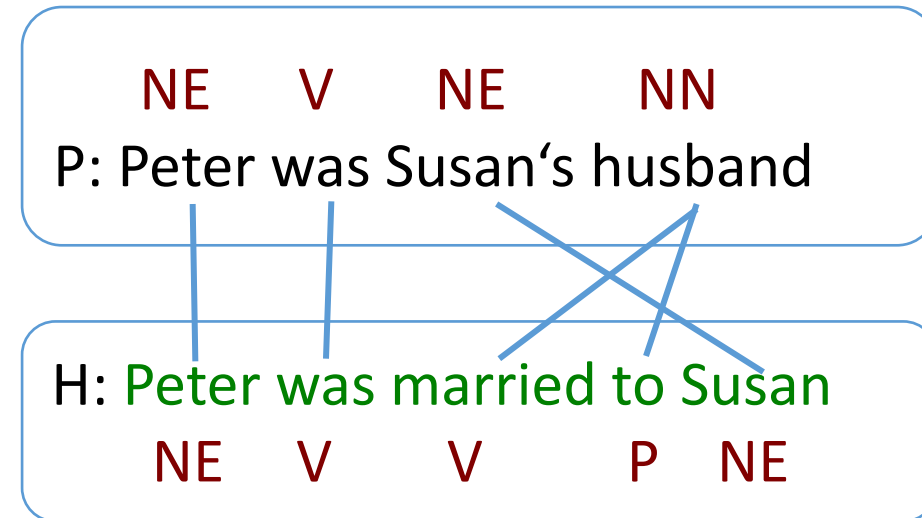


Concrete features

- 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: $5/5 = 100\%$

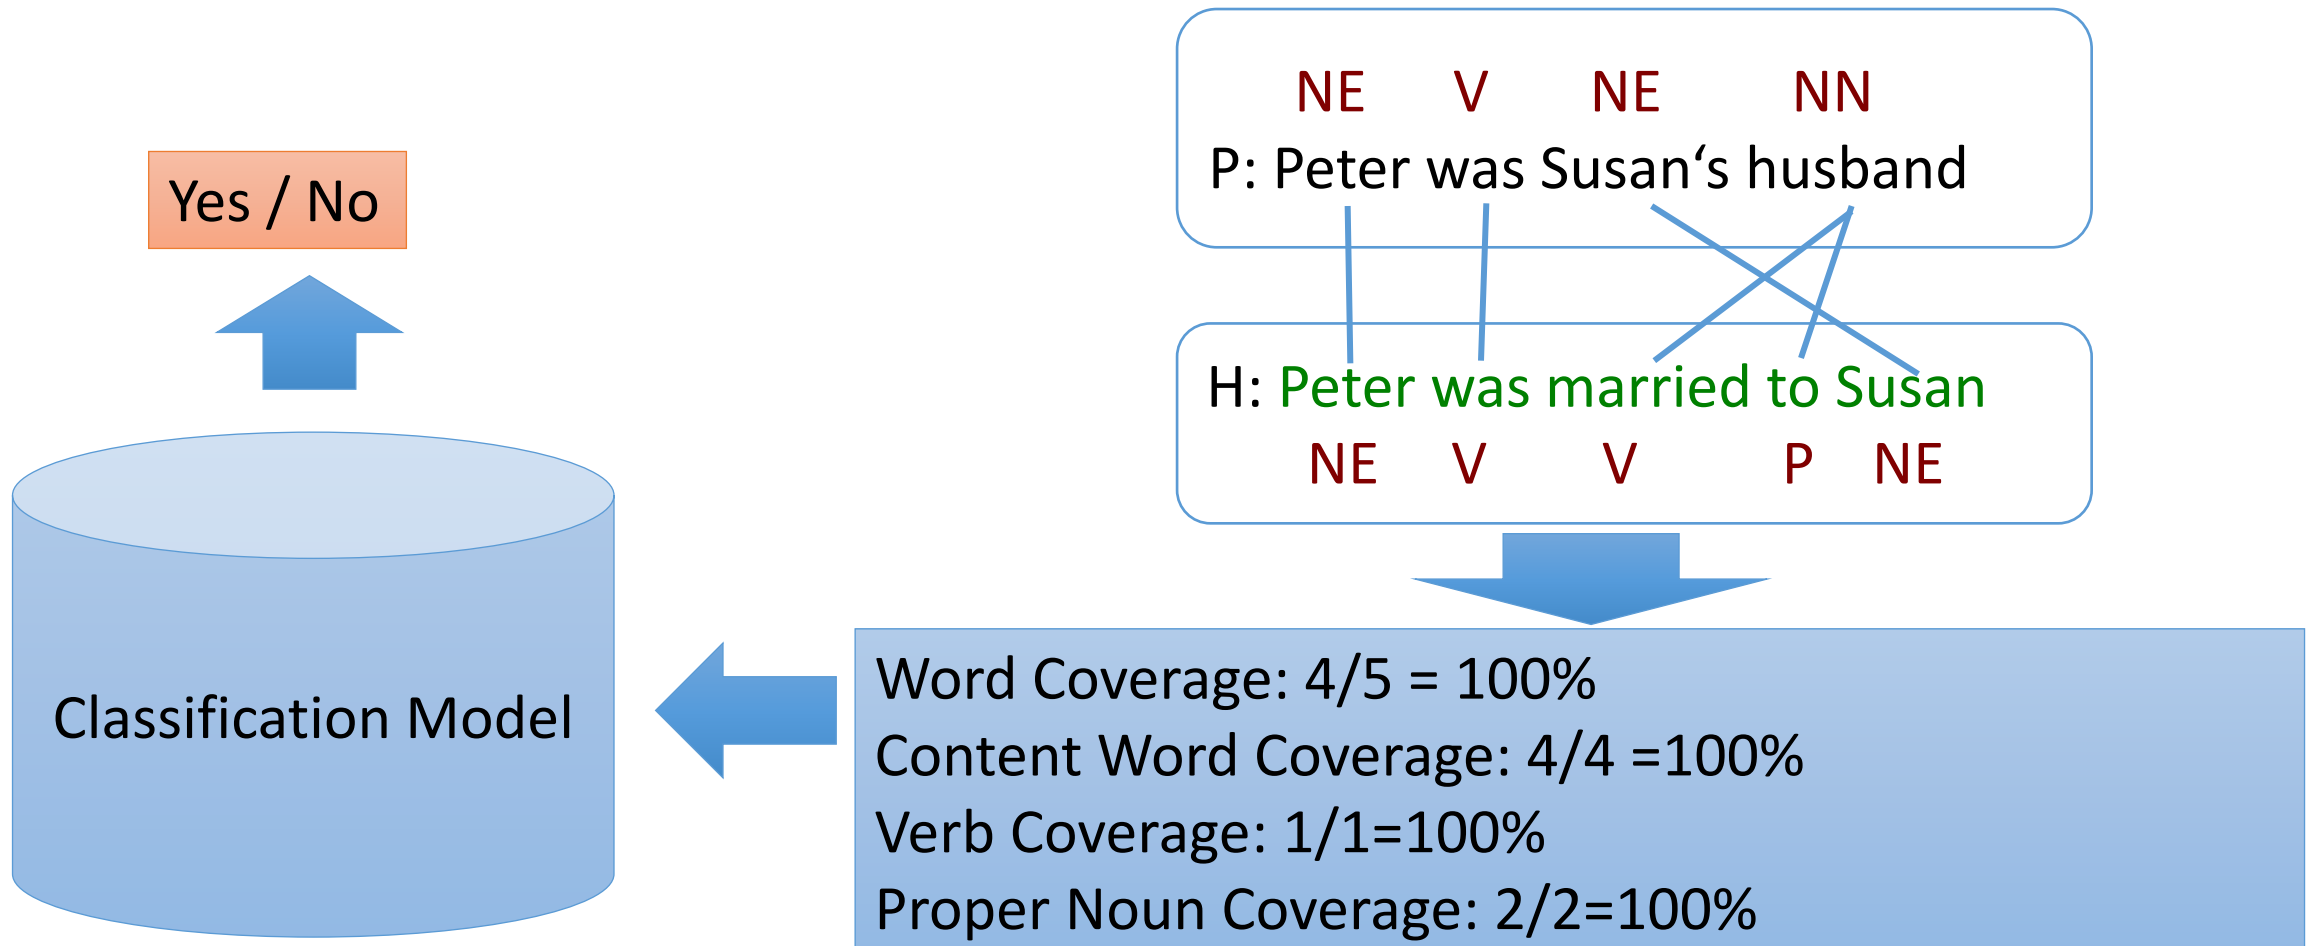
Content Word Coverage: $4/4 = 100\%$

Verb Coverage: $1/1 = 100\%$

Proper Noun Coverage: $2/2 = 100\%$

Alignment-based Entailment: The Algorithmic Level

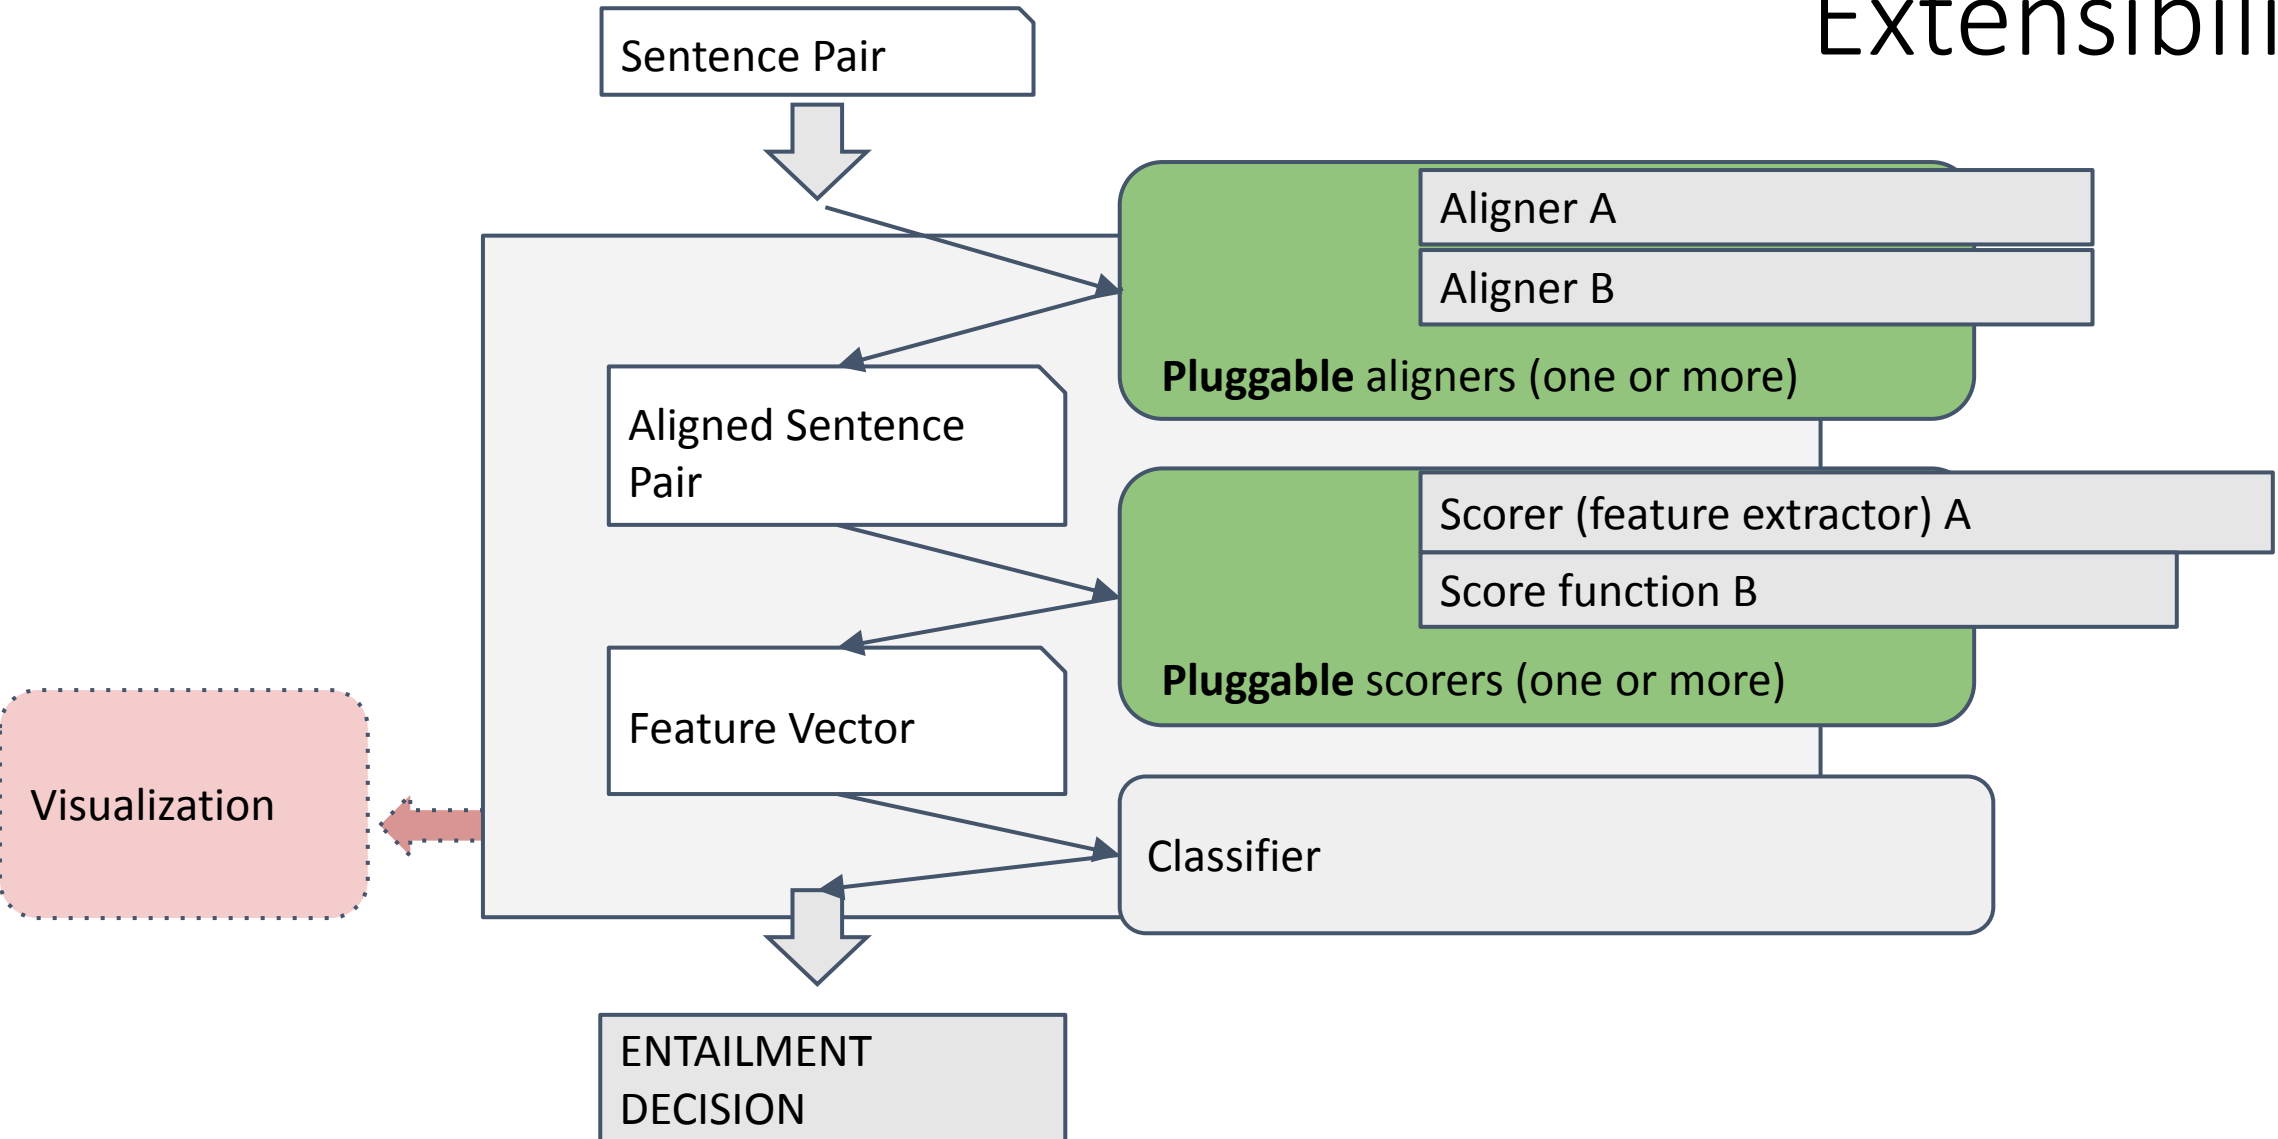
- **Step 4:** Classification (logistic regression, with training examples)



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

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)

- 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 his [Extraterrestrial life](#)

•*Be-complement*

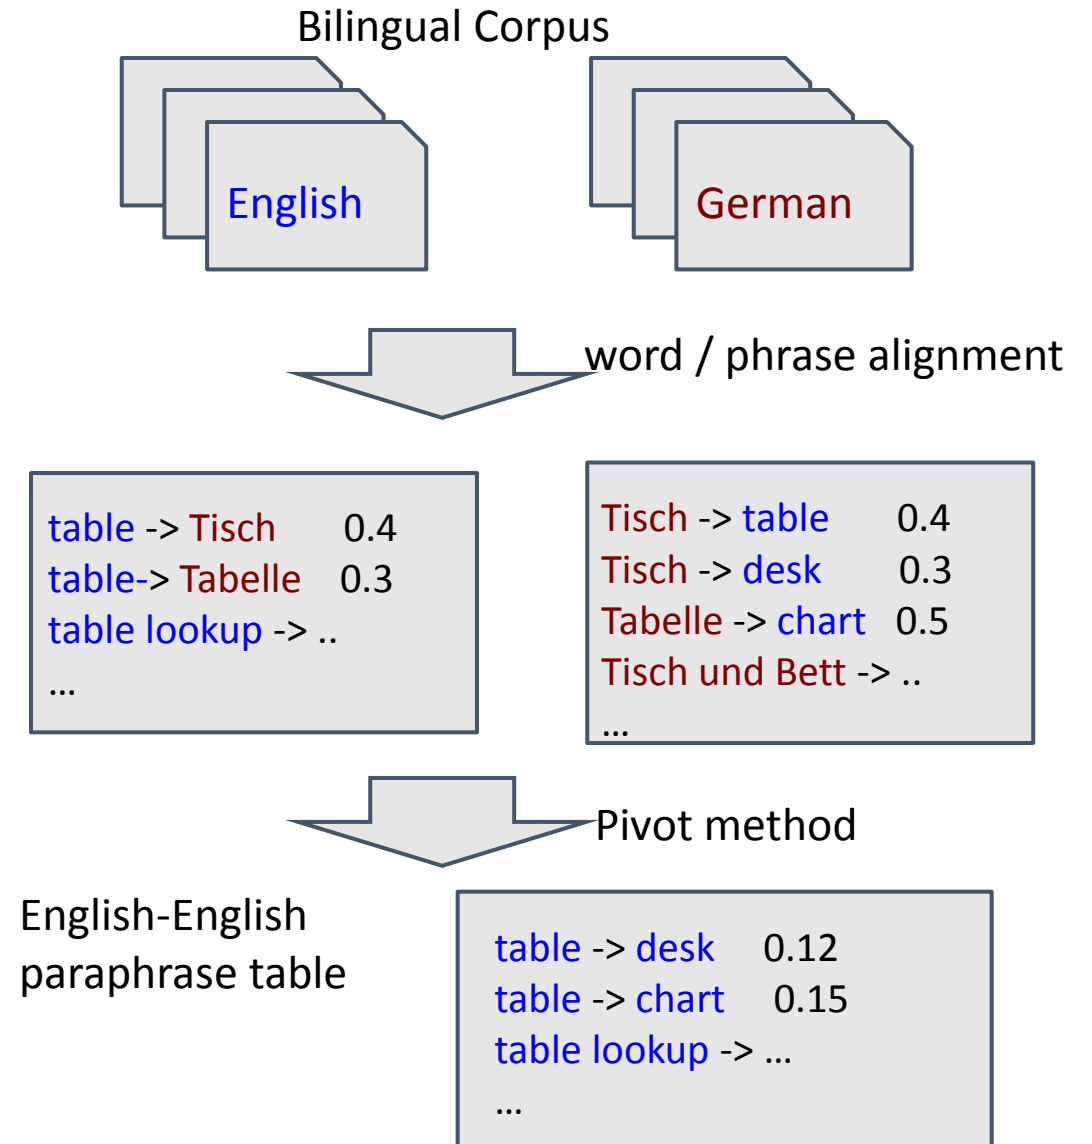
•*Redirect*

•*Parenthesis*

•*Link*

Bilingual-based Paraphrases

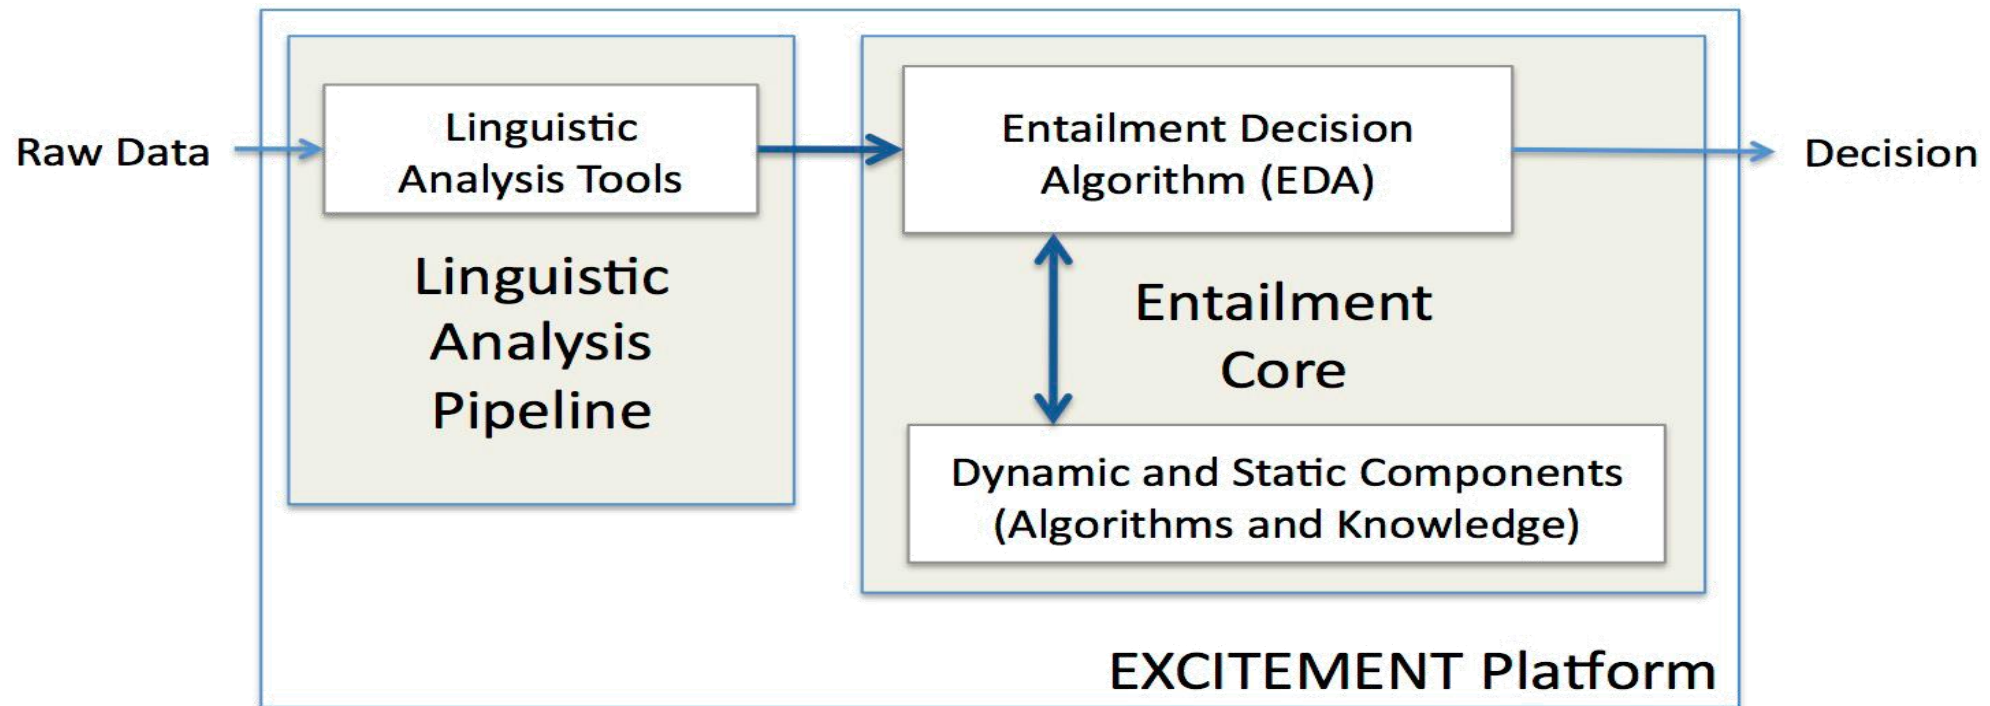
- Intuition: p and p' are paraphrases if both translate into same phrase t (a “pivot”)
- Procedure:
 1. Word- and phrase-align parallel corpus (e.g. English-German)
 2. Extract bilingual translation table
 3. Hop from English to German and back to obtain paraphrase table (plus probability)



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

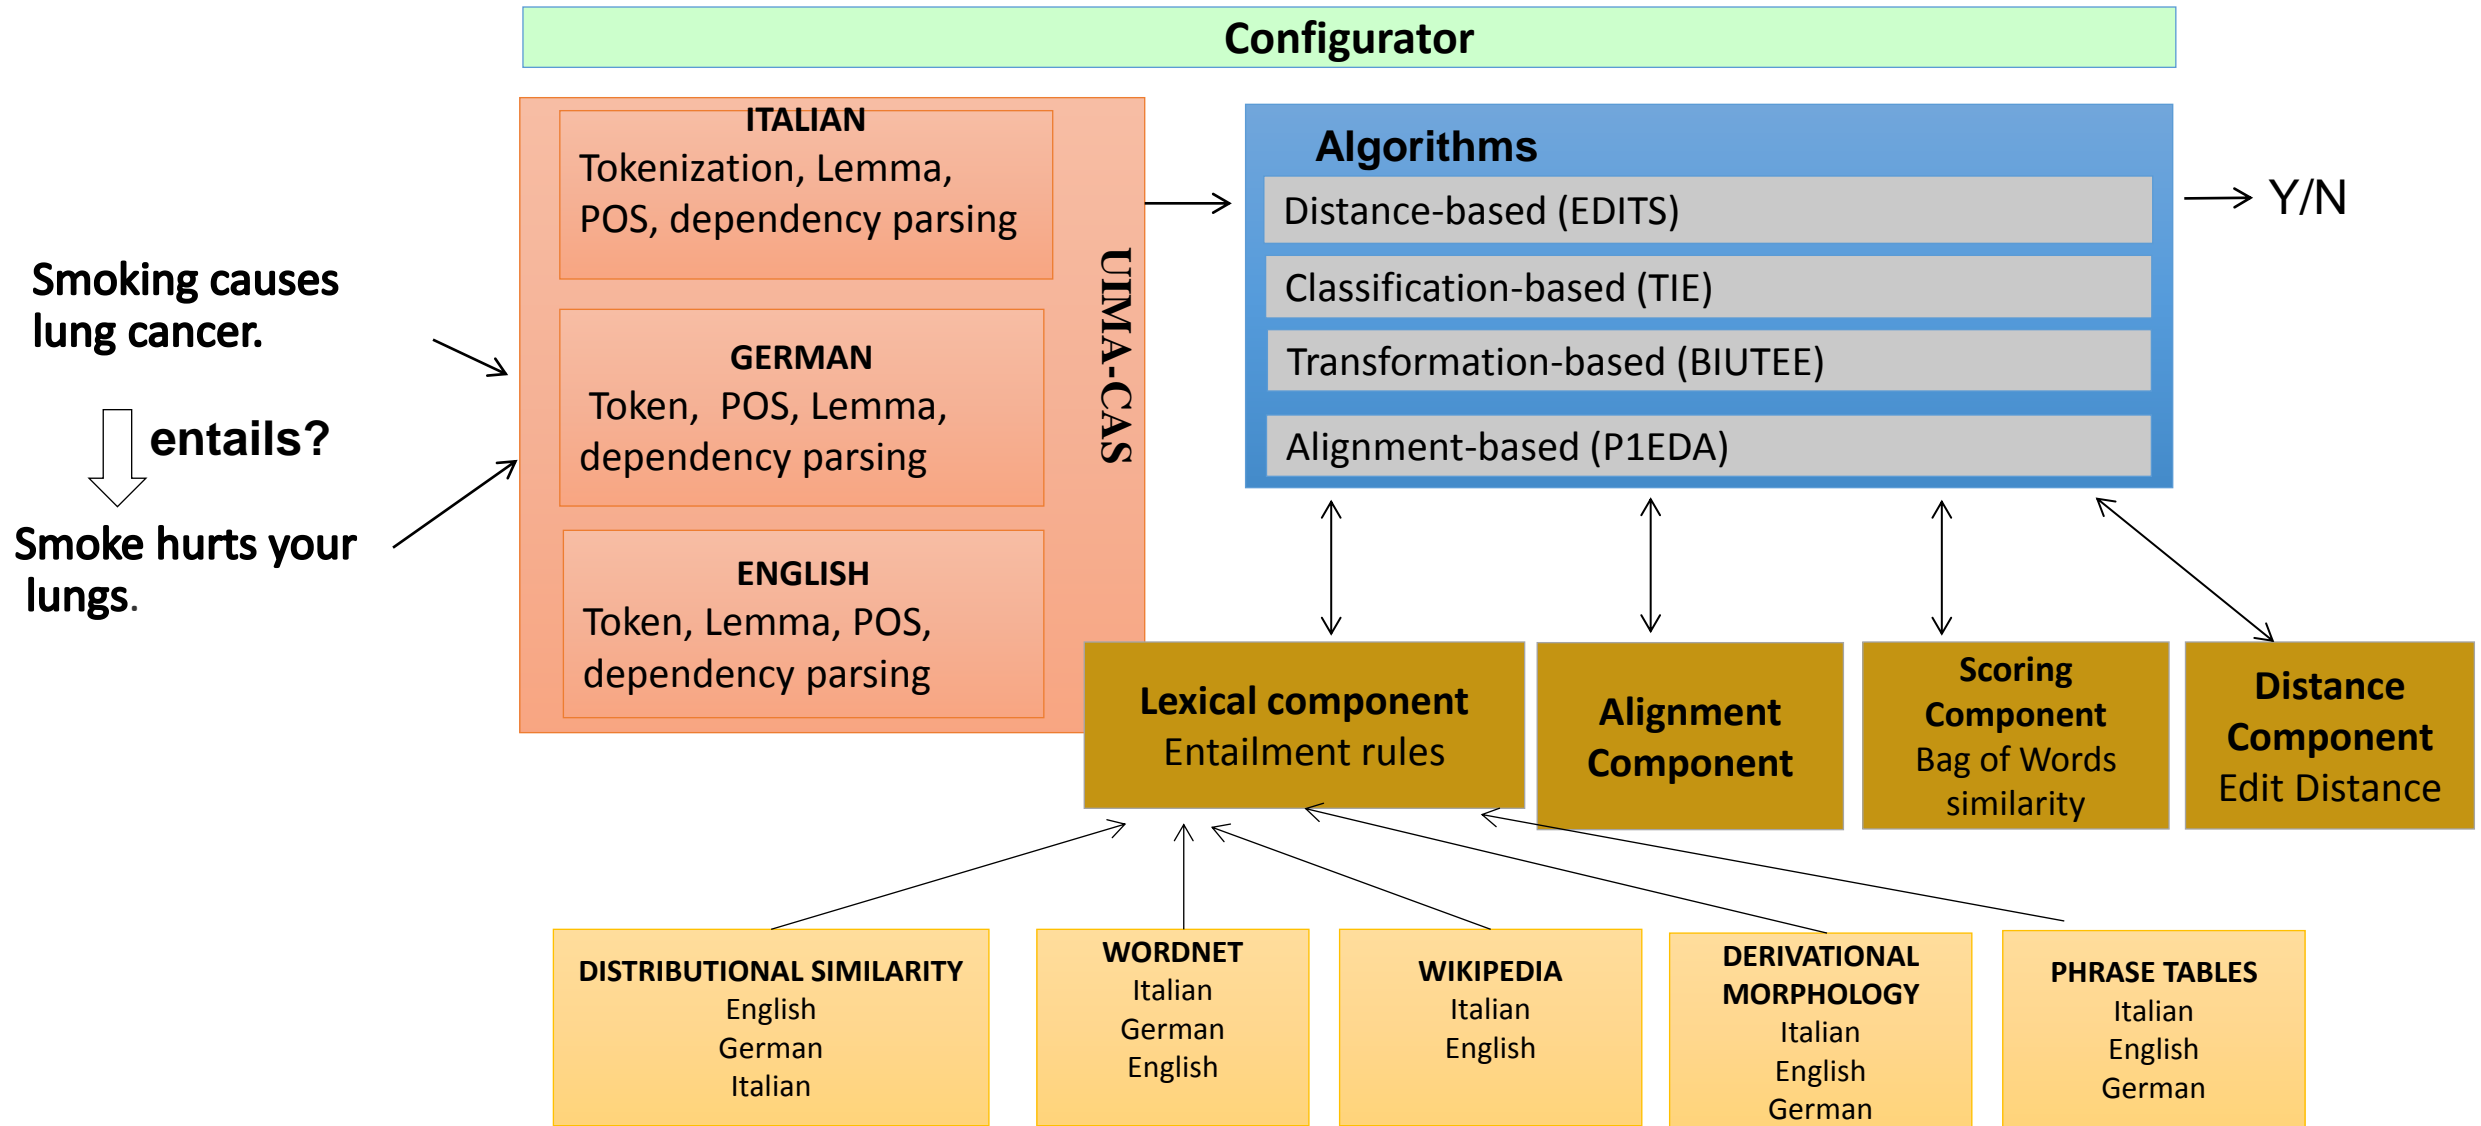


EXCITEMENT EU project: <http://www.excitement-project.eu>

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.githul

Home

Edit New Page

rzanoli edited this page 27 days ago · 11 revisions

The **EXCITEMENT Open Platform (EOP)** is an open source software platform containing state-of-the-art algorithms for recognizing textual entailment relations.

Pages 81

Major changes of release 1.1.4 compared to the previous release 1.1.3:

New Features:

- MaltParser for Italian language

Bug Fixes:

- English MaltParser pipeline made wrong results due to POS tag mismatch.
- Italian TreeTagger missed "canonical" POS tag.
- Util submodule used a wrong version of LAP.

Documentation

- Licence
- Requirements
- Installation
- Quick Start
- Step by Step Tutorial
- Entailment Algorithms
 - BIUTEE
 - EDITS
 - TIE
- 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

EOP
Excitement Open Platform for Recognizing Textual Entailment

// 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.

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

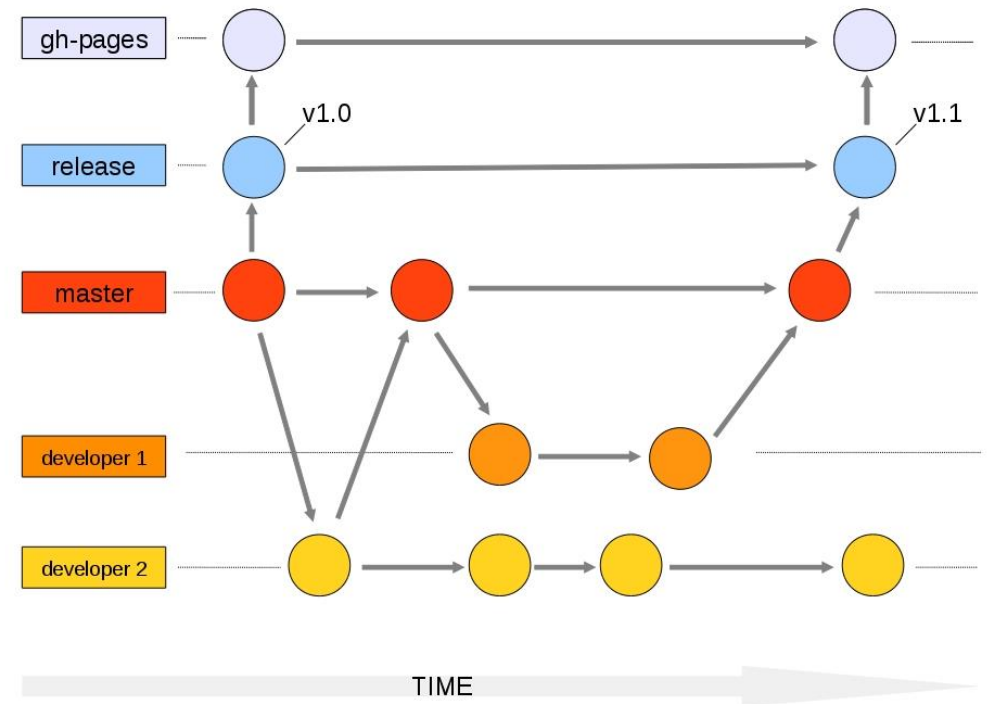
Affiliation: Deutsche Forschungszentrum für Künstliche Intelligenz GmbH (DFKI)

Contact Information: neumann [at] 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 Aligement-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.

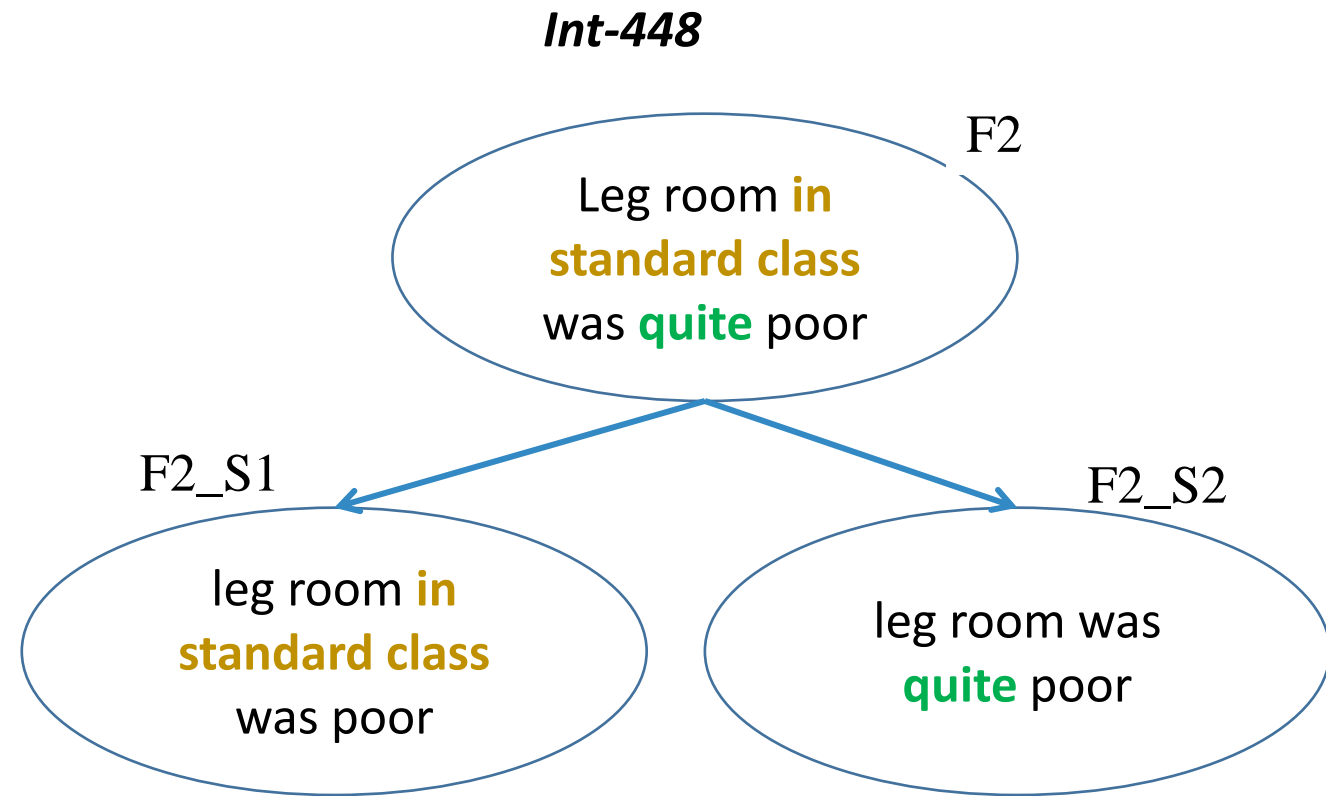
Int-303: My only gripes r **not enough leg room in standard** and I think it would be chic to have refreshments served in carriages , either trolley or trays like in theatres .

Building Fragment Graphs

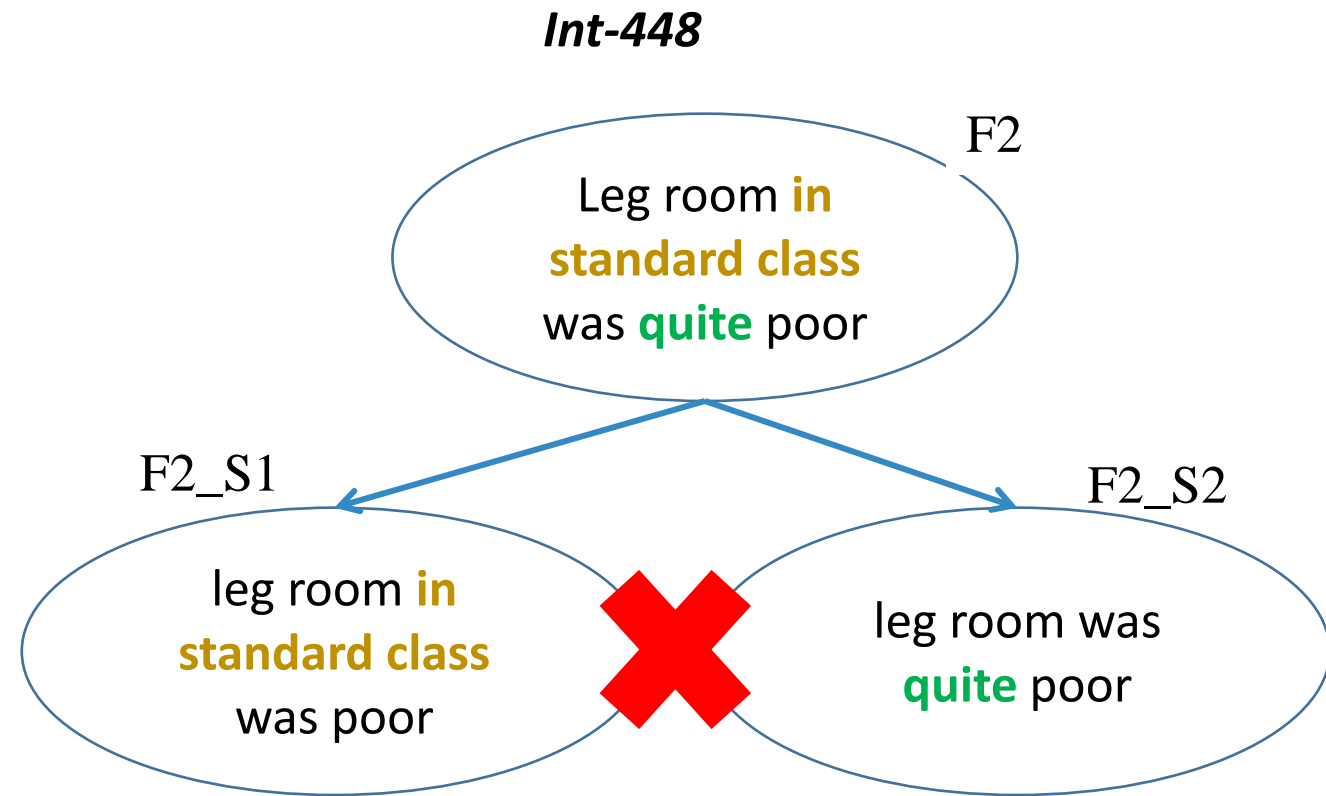
Int-448



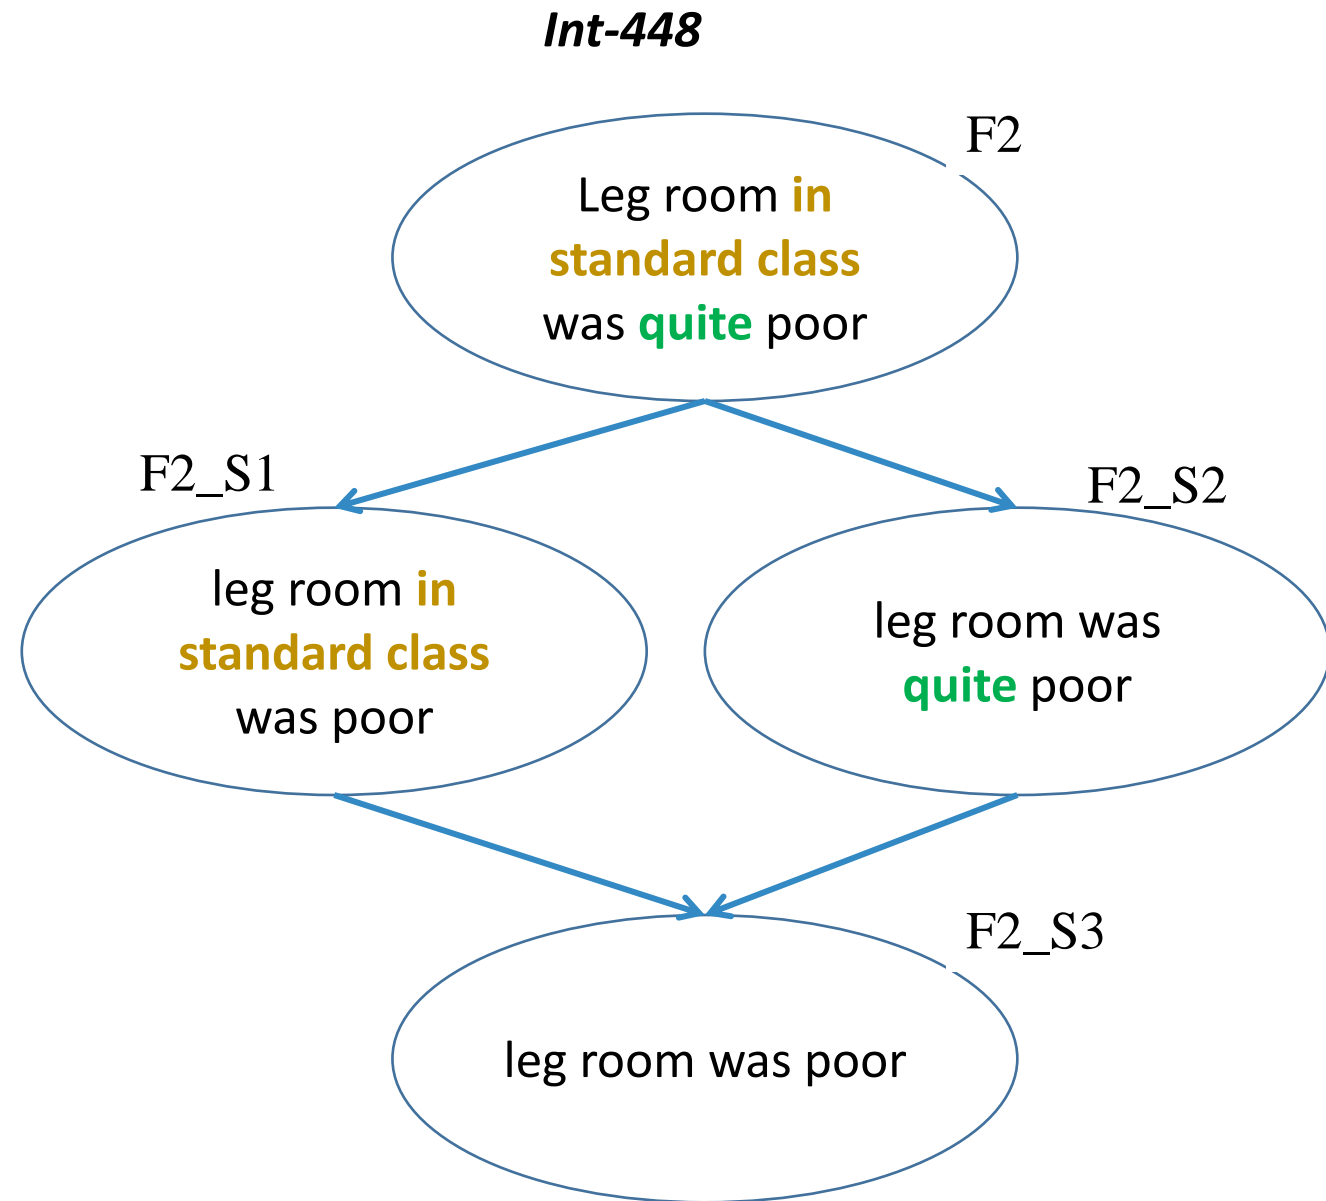
Building Fragment Graphs



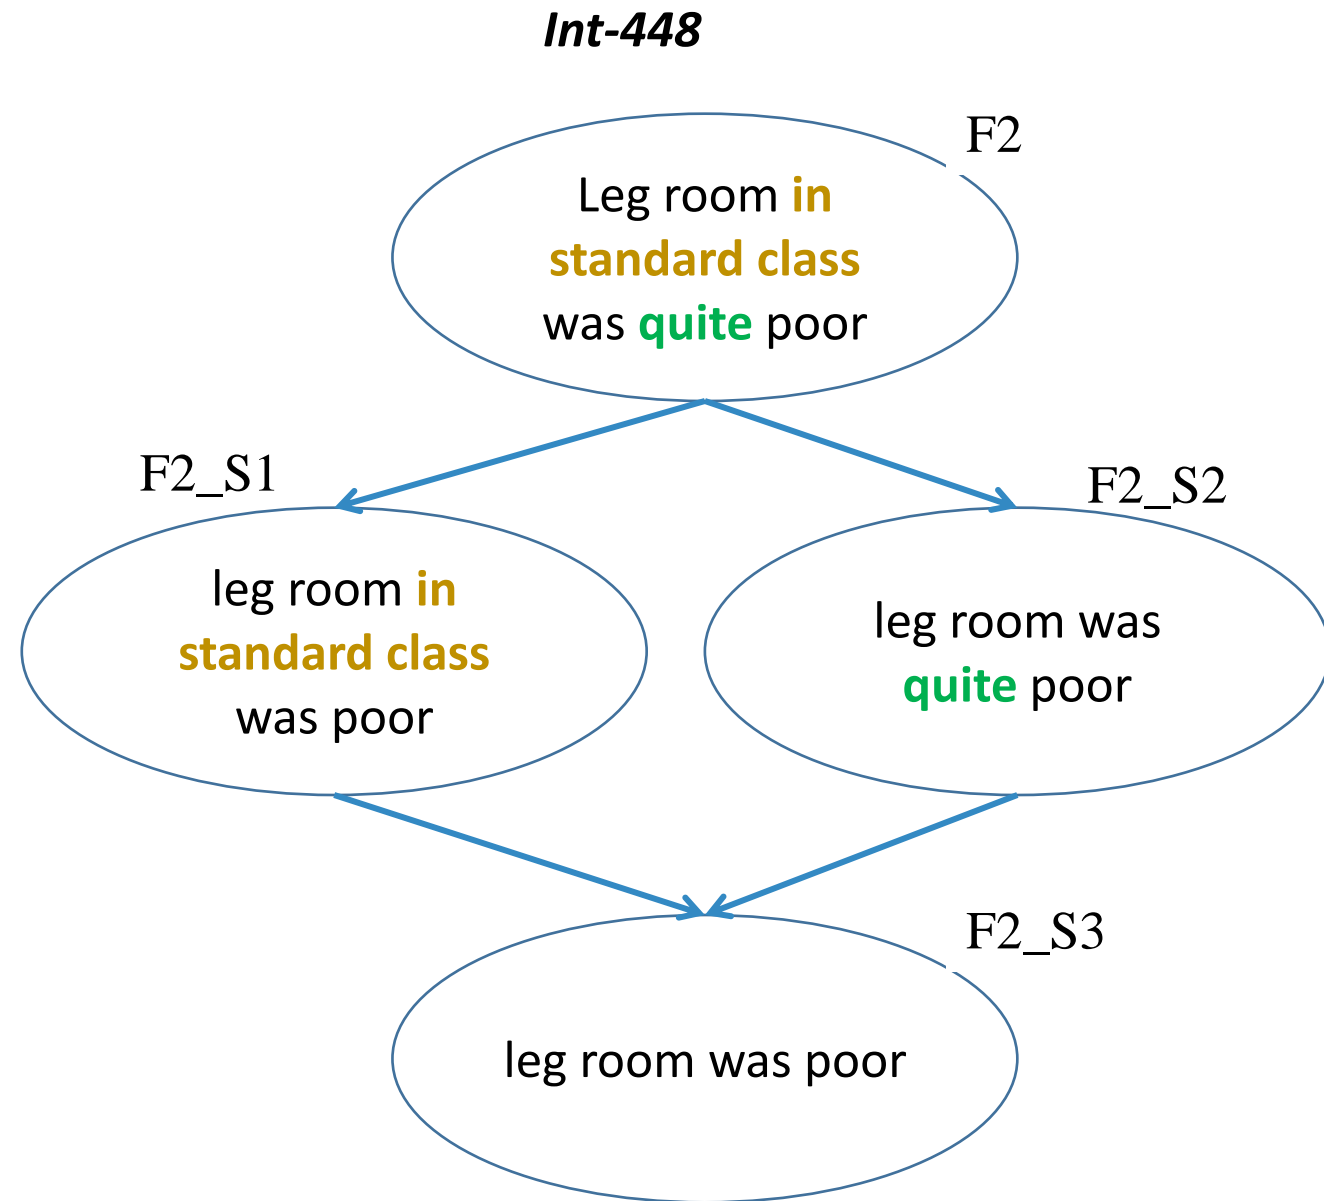
Building Fragment Graphs



Building Fragment Graphs



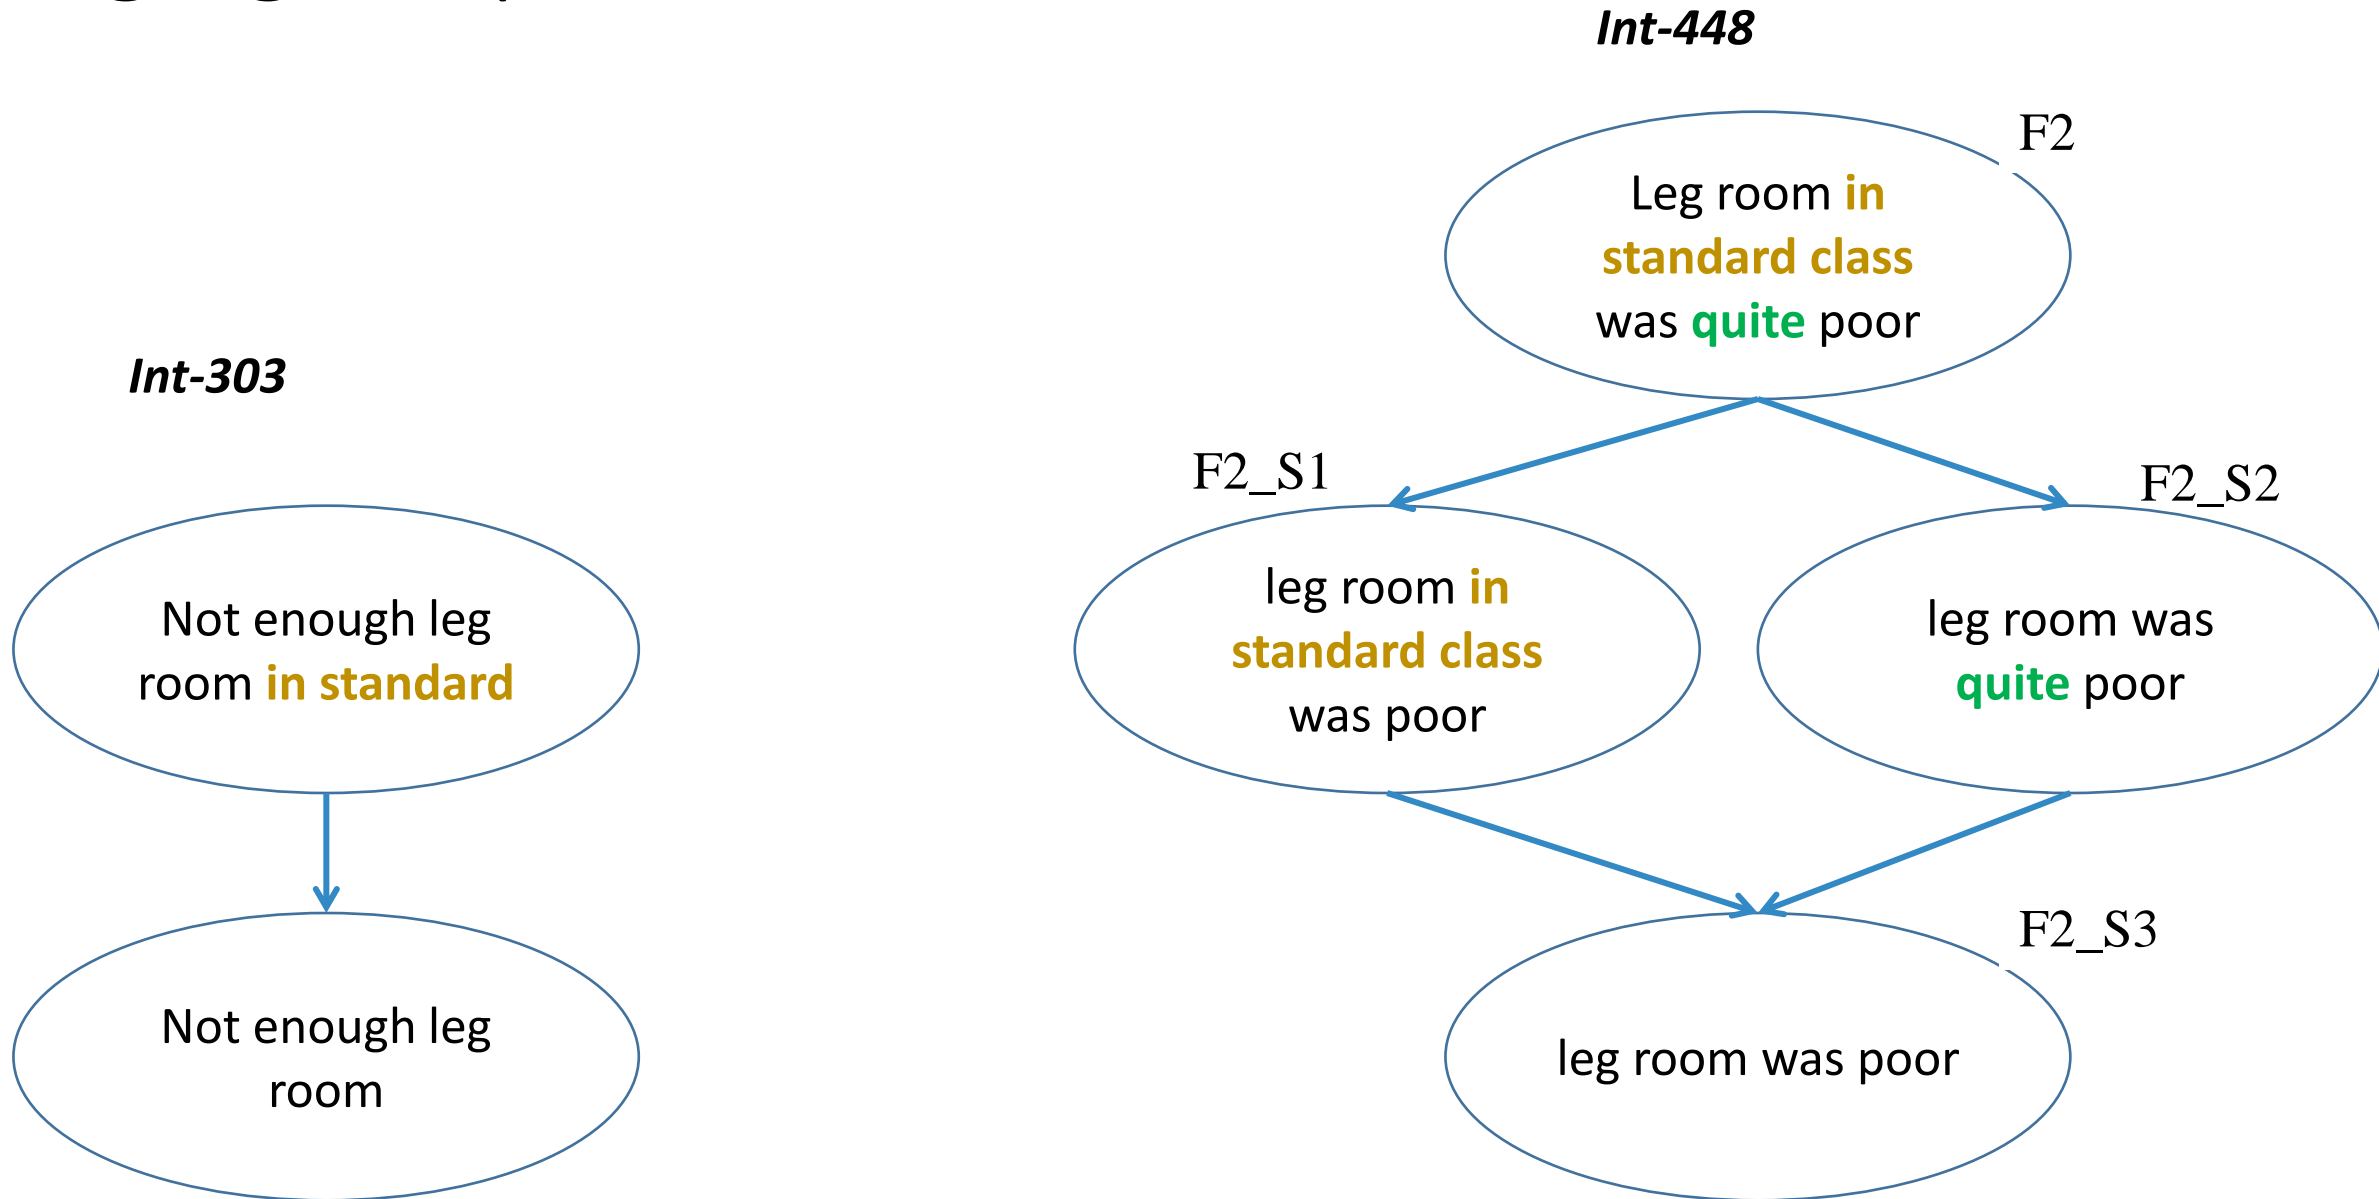
Building Fragment Graphs



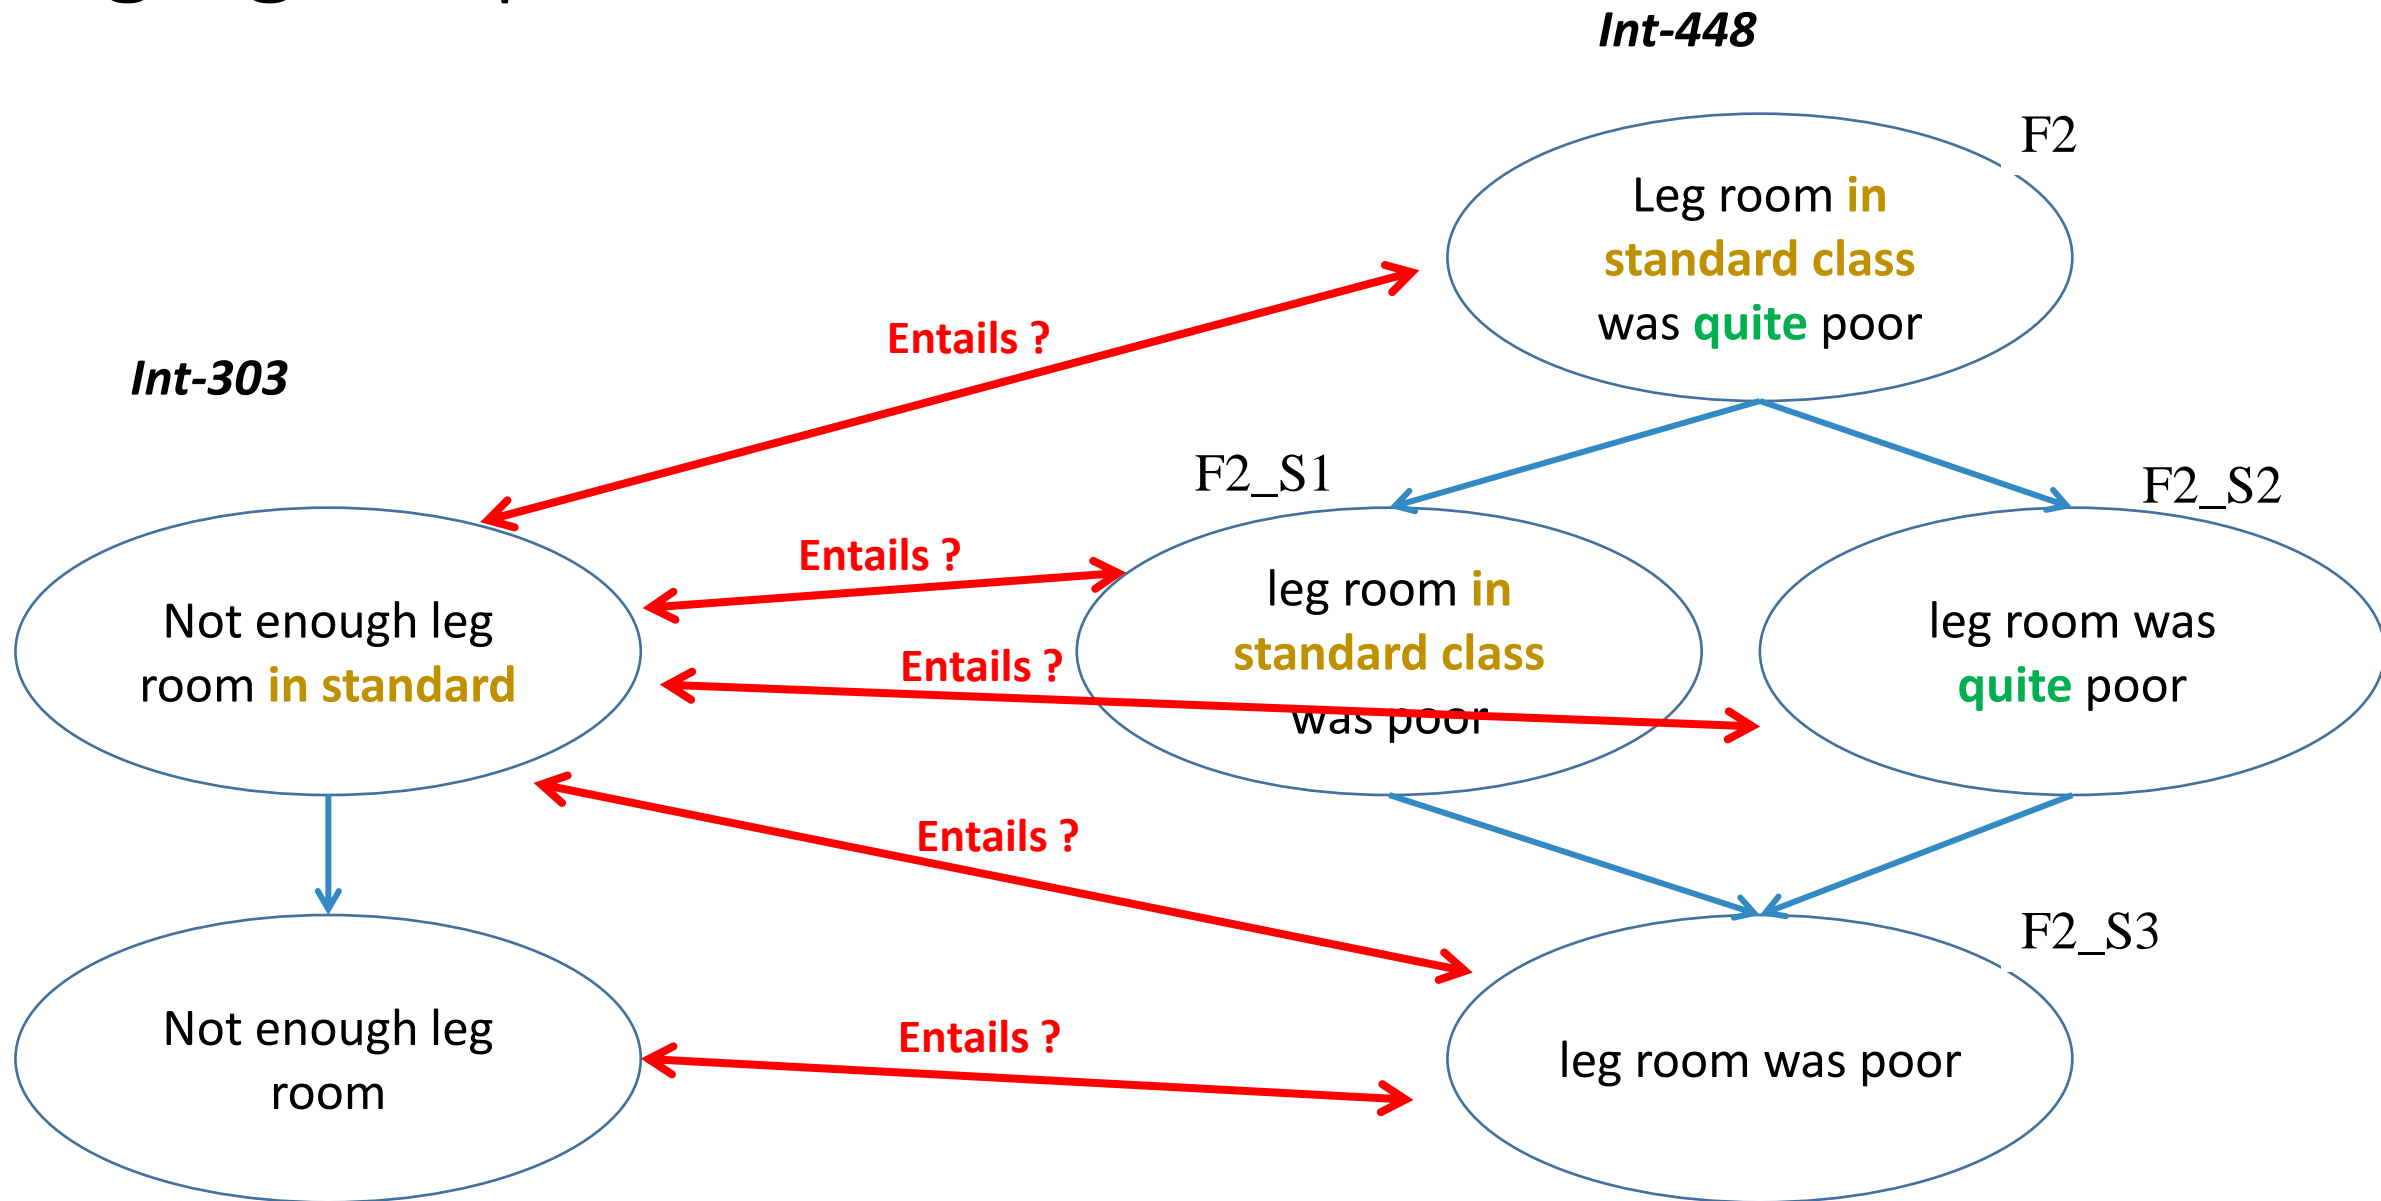
Result: a DAG

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

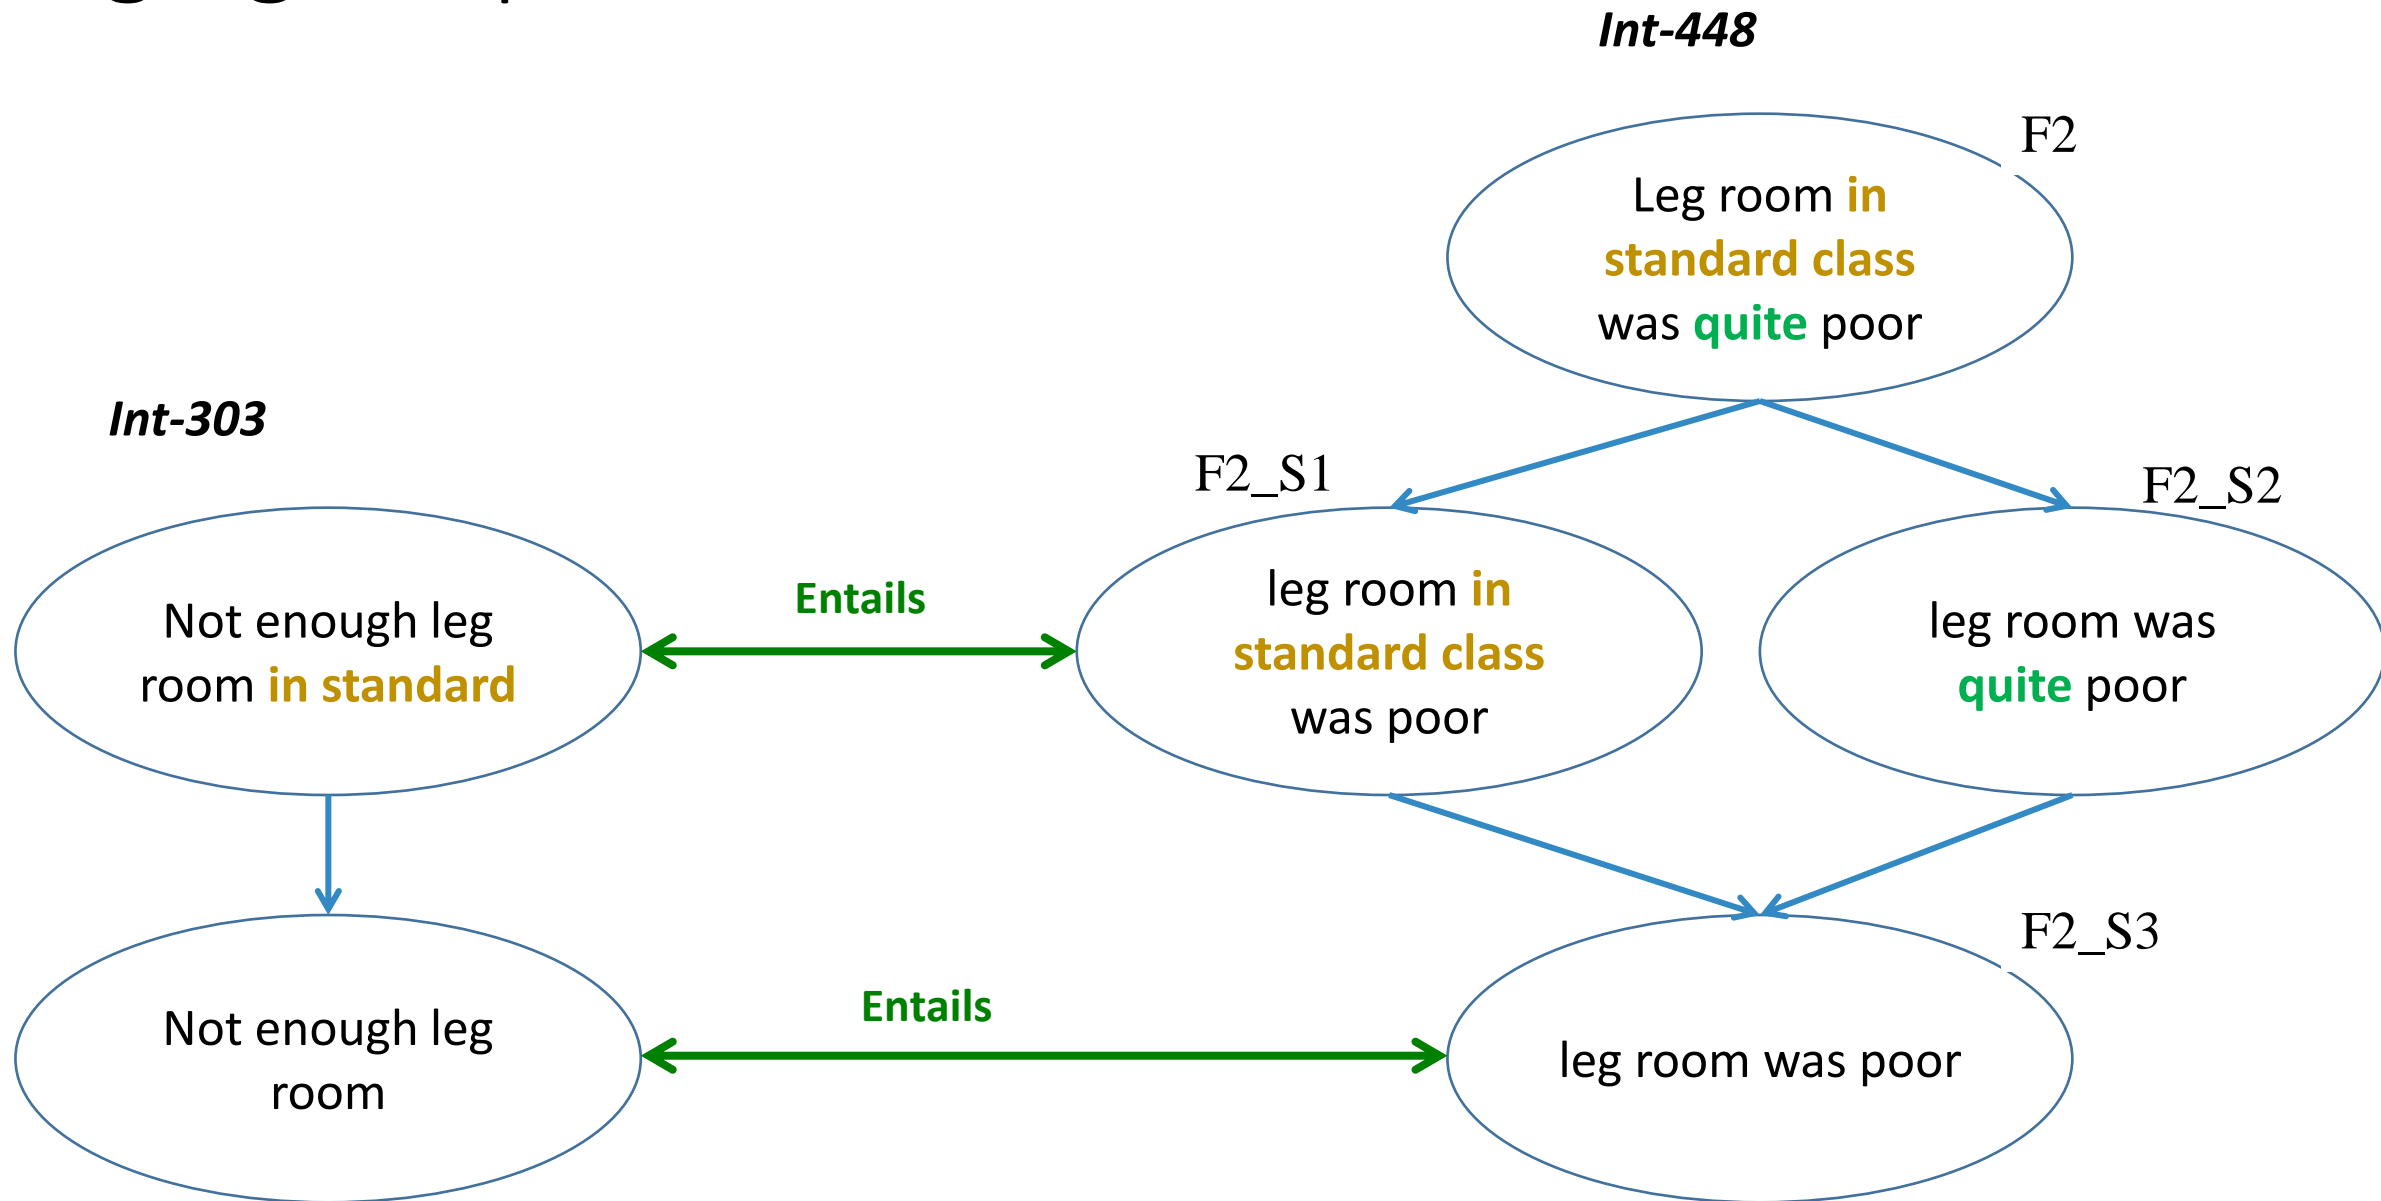
Merging Graphs with the EOP



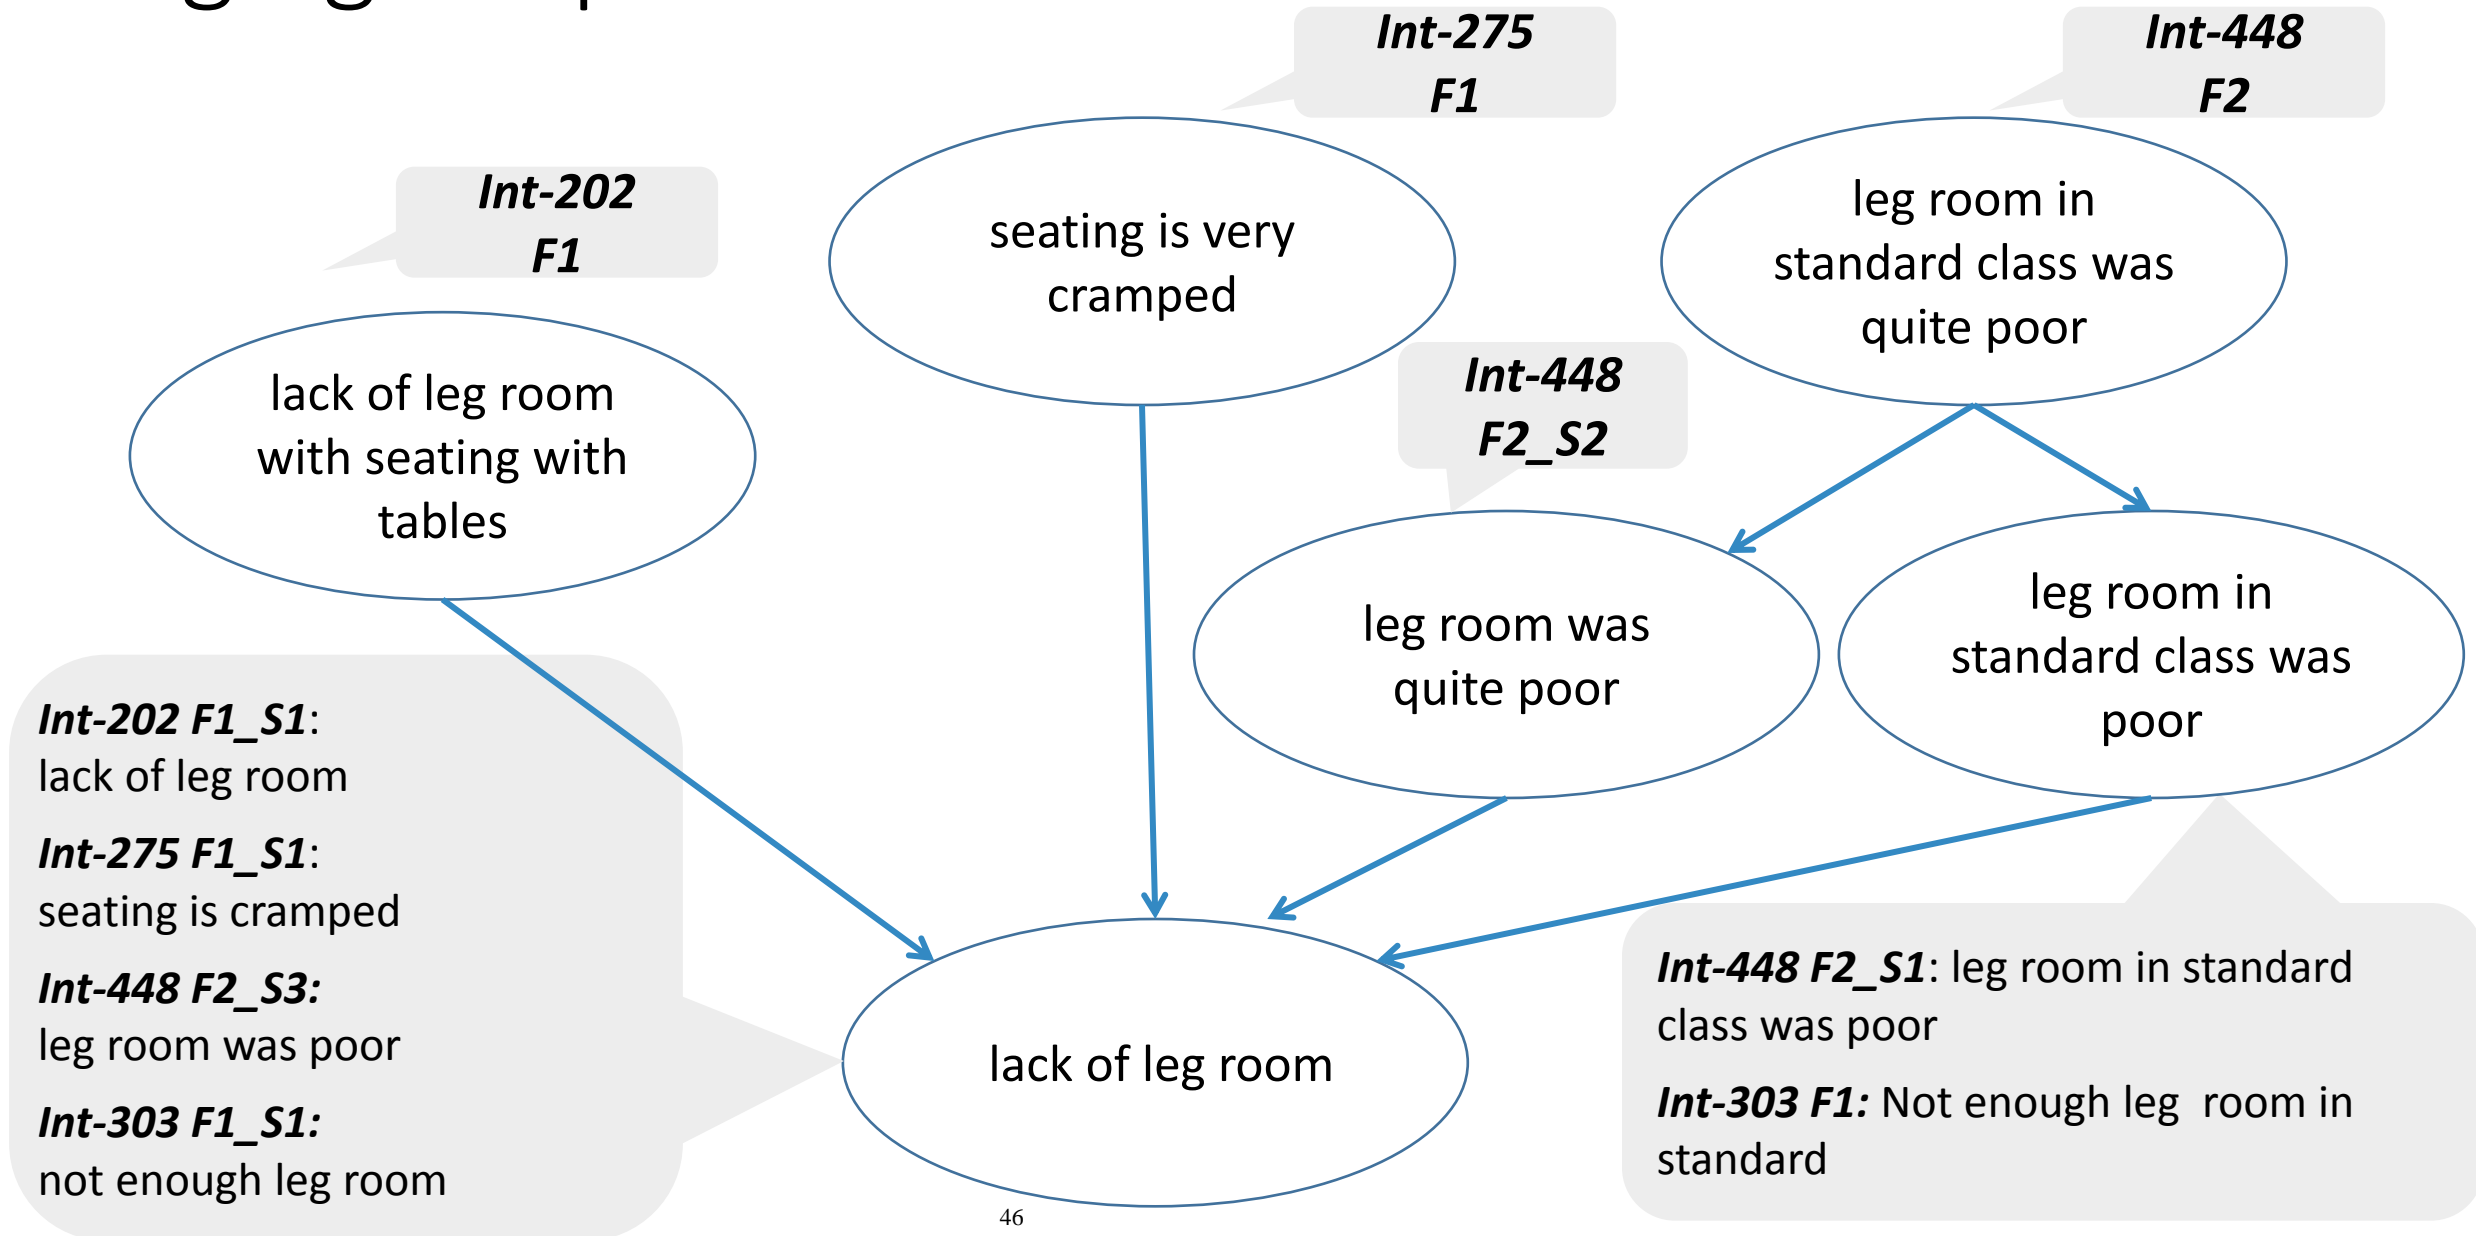
Merging Graphs with the EOP



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