Generating FrameNets of various granularities: The FrameNet Transformer

Josef Ruppenhofer, Jonas Sunde, & Manfred Pinkal

Saarland University

LREC, May 2010

Introduction

- Predicate-argument structure has proven essential for many NLP applications

Introduction

- Predicate-argument structure has proven essential for many NLP applications
- Two prominent resources for modelling predicate-argument structure in English are PropBank (Palmer et al., 2005) and FrameNet (Baker et al., 1998)

Introduction

- Predicate-argument structure has proven essential for many NLP applications
- Two prominent resources for modelling predicate-argument structure in English are PropBank (Palmer et al., 2005) and FrameNet (Baker et al., 1998)
- PropBank maps different syntactic realizations of one lemma to the same predicate-argument structure, using lemma-specific semantic roles

Introduction

- Predicate-argument structure has proven essential for many NLP applications
- Two prominent resources for modelling predicate-argument structure in English are PropBank (Palmer et al., 2005) and FrameNet (Baker et al., 1998)
- PropBank maps different syntactic realizations of one lemma to the same predicate-argument structure, using lemma-specific semantic roles
- FrameNet offers additional structure and detail, making it attractive for information-access tasks

Pros and Cons of Using FrameNet

Pros

Cons

Pros and Cons of Using FrameNet

Pros

- Detail and richness
  - Word senses grouped into Frames
  - Several types of frame relations
  - Parallel to frame relations, FE relations

Cons

Pros and Cons of Using FrameNet

Pros

- Detail and richness
  - Word senses grouped into Frames
  - Several types of frame relations
  - Parallel to frame relations, FE relations

Cons

- Many units are exemplified by relatively few annotated training instances (e.g. Kaisser & Webber 2007).

Pros and Cons of Using FrameNet

Pros

- Detail and richness
  - Word senses grouped into Frames
  - Several types of frame relations
  - Parallel to frame relations, FE relations

Cons

- Many units are exemplified by relatively few annotated training instances (e.g. Kaisser & Webber 2007).
- Distinctions often too fine-grained (Burchardt et al. 2009) to allow robust shallow semantic parsing.

Coarsening FrameNet

- We address the problems of data sparsity and too fine distinctions by coarsening FrameNet with the FN transformer tool

Coarsening FrameNet

- We address the problems of data sparsity and too fine distinctions by coarsening FrameNet with the FN transformer tool
- The tool efficiently generates coarser-grained variants of the FrameNet database.

Coarsening FrameNet

- We address the problems of data sparsity and too fine distinctions by coarsening FrameNet with the FN transformer tool
- The tool efficiently generates coarser-grained variants of the FrameNet database.
  - it reduces the number of word-senses (frames) per lemma
  - it increases the number of annotated sentences per lexical unit and frame.

Coarsening FrameNet

- We address the problems of data sparsity and too fine distinctions by coarsening FrameNet with the FN transformer tool
- The tool efficiently generates coarser-grained variants of the FrameNet database.
  - it reduces the number of word-senses (frames) per lemma
  - it increases the number of annotated sentences per lexical unit and frame.
- we achieve this in two ways
  - merging Frames
  - merging LUs

Coarsening FrameNet

- We address the problems of data sparsity and too fine distinctions by coarsening FrameNet with the FN transformer tool
- The tool efficiently generates coarser-grained variants of the FrameNet database.
  - it reduces the number of word-senses (frames) per lemma
  - it increases the number of annotated sentences per lexical unit and frame.
- we achieve this in two ways
  - merging Frames
  - merging LUs

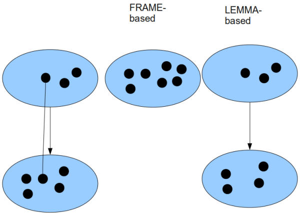

Merging by frame

- The idea is to merge frames in a principled way: by frame-relation

Merging by frame

- The idea is to merge frames in a principled way: by frame-relation
- Merging of senses would result as a side effect

Merging by frame

- The idea is to merge frames in a principled way: by frame-relation
- Merging of senses would result as a side effect
- Frame relations are redirected as needed

Merging by frame

- The idea is to merge frames in a principled way: by frame-relation
- Merging of senses would result as a side effect
- Frame relations are redirected as needed
- Parameters
  - selection of frames that receive annotations
  - selection of frames that disappear
  - stop frames (e.g. Event, Entity,...)

Choosing suitable relations

- Good candidates
  - PERSPECTIVE ON (`Hiring` → `Employment start` ← `Get a job`)
  - SUBFRAME (`Criminal process` → `Arrest`, `Arraignment`, ...)
  - CAUSATIVE OF (`Killing` → `Death`)
  - INCHOATIVE OF (`Death` → `Dead or alive`)
- Less reliable
  - USING (`Communication` → `Volubility`)
  - INHERITANCE (`Transitive action` → `Cause to end`)

Crime scenario original

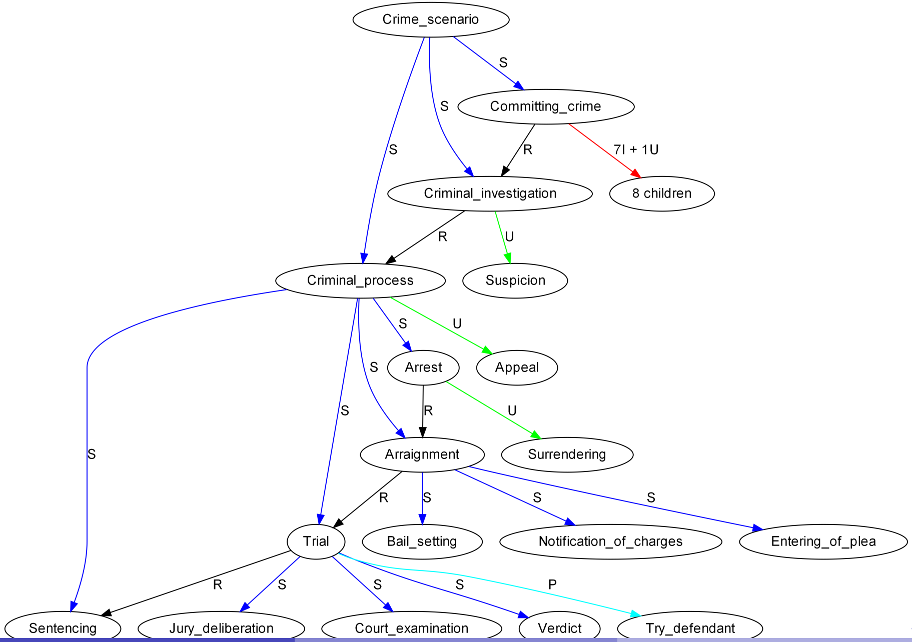

Crime scenario after 1 iteration of frame-based merging

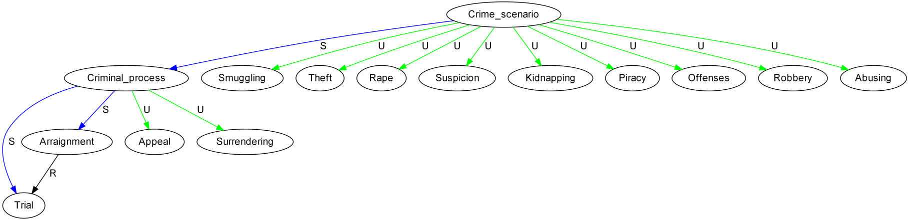

Crime scenario after 2nd iteration of frame-based merging

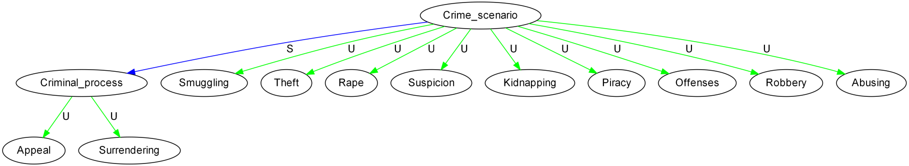

Lemma-based mode

- Merges and migrates related (frame-specific) senses of a particular lemma
- Frame structure remains intact

Lemma-based mode

- Merges and migrates related (frame-specific) senses of a particular lemma
- Frame structure remains intact
- FN release 1.3 has 1316 lemmas that occur in more than one frame.

Lemma-based mode

- Merges and migrates related (frame-specific) senses of a particular lemma
- Frame structure remains intact
- FN release 1.3 has 1316 lemmas that occur in more than one frame.
- Mostly they are involved in polysemy between 2 known senses but in some cases a lemma belongs to 9 different frames.

Lemma-based mode

- Merges and migrates related (frame-specific) senses of a particular lemma
- Frame structure remains intact
- FN release 1.3 has 1316 lemmas that occur in more than one frame.
- Mostly they are involved in polysemy between 2 known senses but in some cases a lemma belongs to 9 different frames.
- These 1316 lemmas have a total of 2587 pairs of senses that could potentially be merged.

Lemma-based mode II

- Two cases
  - one LU's frame is ancestor of the other LU's frame (530 potential pairs to merge)

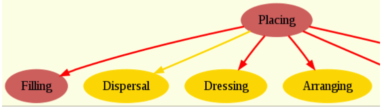

Lemma-based mode II

- Two cases
  - one LU's frame is ancestor of the other LU's frame (530 potential pairs to merge)
  - neither LU's frame is an ancestor for the other: create a new LU in a third frame, reflecting the broader semantic range covered by the combination.

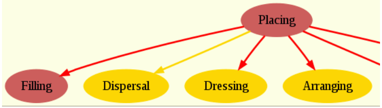

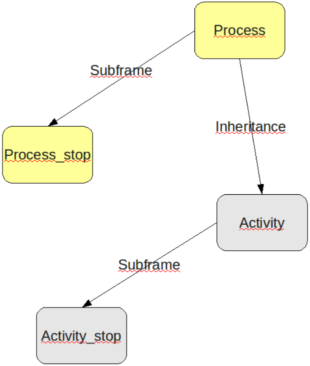

Lemma-based mode II

- Two cases
  - one LU's frame is ancestor of the other LU's frame (530 potential pairs to merge)
  - neither LU's frame is an ancestor for the other: create a new LU in a third frame, reflecting the broader semantic range covered by the combination.
- user selects the types of relations to cross on the path from source to target LUs

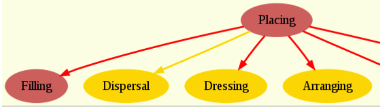

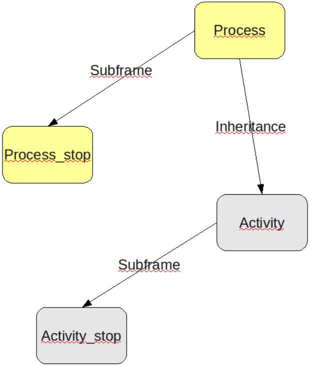

The FN Transformer

- Java 1.6
- no gui, reads user settings from xml file

The FN Transformer

- Java 1.6
- no gui, reads user settings from xml file
- basic settings
  - path to FrameNet data release
  - path to an output directory
  - logfile to be created

The FN Transformer

- Java 1.6
- no gui, reads user settings from xml file
- basic settings
  - path to FrameNet data release
  - path to an output directory
  - logfile to be created
- output is a format-compliant FrameNet release (xml files)

The FN Transformer

- Java 1.6
- no gui, reads user settings from xml file
- basic settings
  - path to FrameNet data release
  - path to an output directory
  - logfile to be created
- output is a format-compliant FrameNet release (xml files)
- for relation-based merger, output also includes .dot files that can be used for “visual diff” (via the open-source GraphViz software)

The FN Transformer

- Java 1.6
- no gui, reads user settings from xml file
- basic settings
  - path to FrameNet data release
  - path to an output directory
  - logfile to be created
- output is a format-compliant FrameNet release (xml files)
- for relation-based merger, output also includes .dot files that can be used for "visual diff" (via the open-source GraphViz software)
- in addition to the two automatic modes, there is a manual mode

Evaluation

- A baseline evaluation consists in confirming that we do obtain the expected improved **accuracy** of frame-semantic parsers trained on the modified data.
- In a further step, we perform a task-based evaluation to check whether we improve parsing accuracy at the cost of losing **relevant information**.

Parsing accuracy: setup

- Compare the performance of the Shalmaneser semantic parser (Erk & Padó 2006) in two settings:
  - Baseline: FrameNet release 1.3.
  - Coarsened: modified FrameNets created by our transformer
- Data: subset of lemmas that were affected by the transformation
- 10-fold cross-validation setting
  - frame assignment
  - argument recognition
  - argument labeling

Parsing accuracy

	task	cum.	task	cum.
	FN1.3		**FN1.3R**	
Frame assignment	0.94	0.94	0.94	0.94
Argument recognition	0.69	0.64	0.69	0.65
Argument labeling	0.71	0.46	**0.75**	**0.49**

Parsing accuracy

	task	cum.	task	cum.
	FN1.3		**FN1.3R**	
Frame assignment	0.94	0.94	0.94	0.94
Argument recognition	0.69	0.64	0.69	0.65
Argument labeling	0.71	0.46	**0.75**	**0.49**

	FN1.3		**FN1.3LU**	
Frame assignment	0.89	0.89	**0.94**	**0.94**
Argument recognition	0.69	0.62	0.66	0.62
Argument labeling	0.74	0.46	0.72	0.44

Table: Performance of Shalmaneser on FN release 1.3 and on transformations (10-fold cross-validation)

Preservation of relevant information - RTE

- Ensure that better parser performance is not achieved at the cost of losing relevant information
- Evaluate our coarsened FrameNet versions in the context of the entailment recognition (RTE) task
- Entailment recognition is the task of determining whether a text T entails a hypothesis H in a common sense way.

Preservation of relevant information - RTE

- Ensure that better parser performance is not achieved at the cost of losing relevant information
- Evaluate our coarsened FrameNet versions in the context of the entailment recognition (RTE) task
- Entailment recognition is the task of determining whether a text T entails a hypothesis H in a common sense way.

Preservation of relevant information - RTE

- Ensure that better parser performance is not achieved at the cost of losing relevant information
- Evaluate our coarsened FrameNet versions in the context of the entailment recognition (RTE) task
- Entailment recognition is the task of determining whether a text T entails a hypothesis H in a common sense way.

(3) T: An avalanche has struck a popular skiing resort in Austria, killing at least 11 people.
H: Humans died in an avalanche.

Frame semantic information in the RTE task

- Techniques for judging entailment include measuring lexical overlap, shallow syntactic parsing, and the use of WordNet relations

Frame semantic information in the RTE task

- Techniques for judging entailment include measuring lexical overlap, shallow syntactic parsing, and the use of WordNet relations
- Another kind of approach consists in using shallow semantic representations that abstract away from semantically irrelevant variations

(5) T: An avalanche has struck a popular skiing resort in Austria, killing at least 11 people.
H: Humans died in an avalanche.

Assessing the contribution of frame semantics to RTE

- Burchardt et al 2009 performed an experiment on the gold standard data of the FATE corpus (Burchardt and Pennacchiotti 2008)

Assessing the contribution of frame semantics to RTE

- Burchardt et al 2009 performed an experiment on the gold standard data of the FATE corpus (Burchardt and Pennacchiotti 2008)
- FATE corpus: manual frame semantic annotations for the 800 entailment pairs of RTE-2 ; 4490 frame instances annotated.

Assessing the contribution of frame semantics to RTE

- Burchardt et al 2009 performed an experiment on the gold standard data of the FATE corpus (Burchardt and Pennacchiotti 2008)
- FATE corpus: manual frame semantic annotations for the 800 entailment pairs of RTE-2 ; 4490 frame instances annotated.
- **Key assumption**: the more of the semantics of the hypothesis can be embedded into the text, the more likely it is that an entailment relation holds between text and hypothesis

Assessing the contribution of frame semantics to RTE

- Burchardt et al 2009 performed an experiment on the gold standard data of the FATE corpus (Burchardt and Pennacchiotti 2008)
- FATE corpus: manual frame semantic annotations for the 800 entailment pairs of RTE-2 ; 4490 frame instances annotated.
- **Key assumption**: the more of the semantics of the hypothesis can be embedded into the text, the more likely it is that an entailment relation holds between text and hypothesis
- Extracting frame-based statistical information from the positive and negative examples of the annotated corpus, respectively, and measuring the overlap of frame structures between text and hypothesis in an entailment pair.

(9) T: An avalanche has struck a popular skiing resort in Austria, killing at least 11 people.
H: Humans died in an avalanche.

Frame label overlap

	Positive pairs	Negative pairs	Difference
FN1.3	0.5711	0.4585	0.1126
FN1.3R	0.5913	0.4845	0.1068
FN1.3LU	0.5323	0.4348	0.0975

Table: Average frame label overlap on entailment pairs in three versions of the Fate corpus

Conclusion

- We presented a tool for semi-automatically deriving customized but format-compliant versions of the FrameNet database based on frame and frame element relations.

Conclusion

- We presented a tool for semi-automatically deriving customized but format-compliant versions of the FrameNet database based on frame and frame element relations.
- In baseline evaluations, we found that coarsening FrameNet yields slightly better parsing accuracy and does not cause the loss of information for the RTE task

Conclusion

- We presented a tool for semi-automatically deriving customized but format-compliant versions of the FrameNet database based on frame and frame element relations.
- In baseline evaluations, we found that coarsening FrameNet yields slightly better parsing accuracy and does not cause the loss of information for the RTE task
- Allows users to produce FrameNet versions whose granularity is suitable for their particular applications.

Conclusion

- We presented a tool for semi-automatically deriving customized but format-compliant versions of the FrameNet database based on frame and frame element relations.
- In baseline evaluations, we found that coarsening FrameNet yields slightly better parsing accuracy and does not cause the loss of information for the RTE task
- Allows users to produce FrameNet versions whose granularity is suitable for their particular applications.
- Additional experiments needed to assess whether the individual gains of the two modes of transformation can be combined and what the best settings are for each of them.

Visual diff

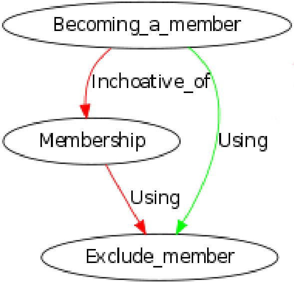

Becoming_a_member
Inchoative_of
Membership
Using
Using
Exclude_member