Relation Extraction for the Food Domain without Labeled Training Data – Is Distant Supervision the Best Solution?

PolTAL 2014

By Melanie Reiplinger, Michael Wiegand,

and Dietrich Klakow

Spoken Language Systems

Saarland University, Germany

Outline of Talk

- Introduction
- Tasks & Data
- Methods and Experiments of Producing Training Data
- Conclusions

Outline of Talk

- Introduction
- Tasks & Data
- Methods and Experiments of Producing Training Data
- Conclusions

Introduction

- Classify (food) relations, using a Distant Supervision approach (DS).
- Optimize training data selection/representation.
- Investigate degrees of freedom in classifier design:
  - knowledge used,
  - processing levels.
- Compare DS against simplistic, rule-based approach.

Outline of Talk

- Introduction
- Tasks & Data
- Methods and Experiments of Producing Training Data
- Conclusions

The Relation Types (I)

- ***SuitsTo(<foodItem>, <foodItem>)***
- Definition:

  Describes food items that are typically consumed together.
- Example: <*hamburger, fries*>

The Relation Types (II)

- ***SubstitutedBy(<foodItem>, <foodItem>)***
- Definition:

  Lists food items that are fairly similar and can be used in similar situations.
- Example: <*butter, margarine*>

Relevance of Relation Types Beyond the Food Domain

Relevance of Relation Types Beyond the Food Domain

- The relation type ***SuitsTo*** is relevant to many other domains.

Relevance of Relation Types Beyond the Food Domain

- The relation type ***SuitsTo*** is relevant to many other domains.

Relevance of Relation Types Beyond the Food Domain

- The relation type ***SuitsTo*** is relevant to many other domains.
- The same holds for the relation type ***SubstitutedBy***.

Possible Applications

Virtual Customer Advice

- Assist a customer in a supermarket.
- Support with knowledge not available via other available modes of information (e.g. packaging).

Domain-Specific Corpus

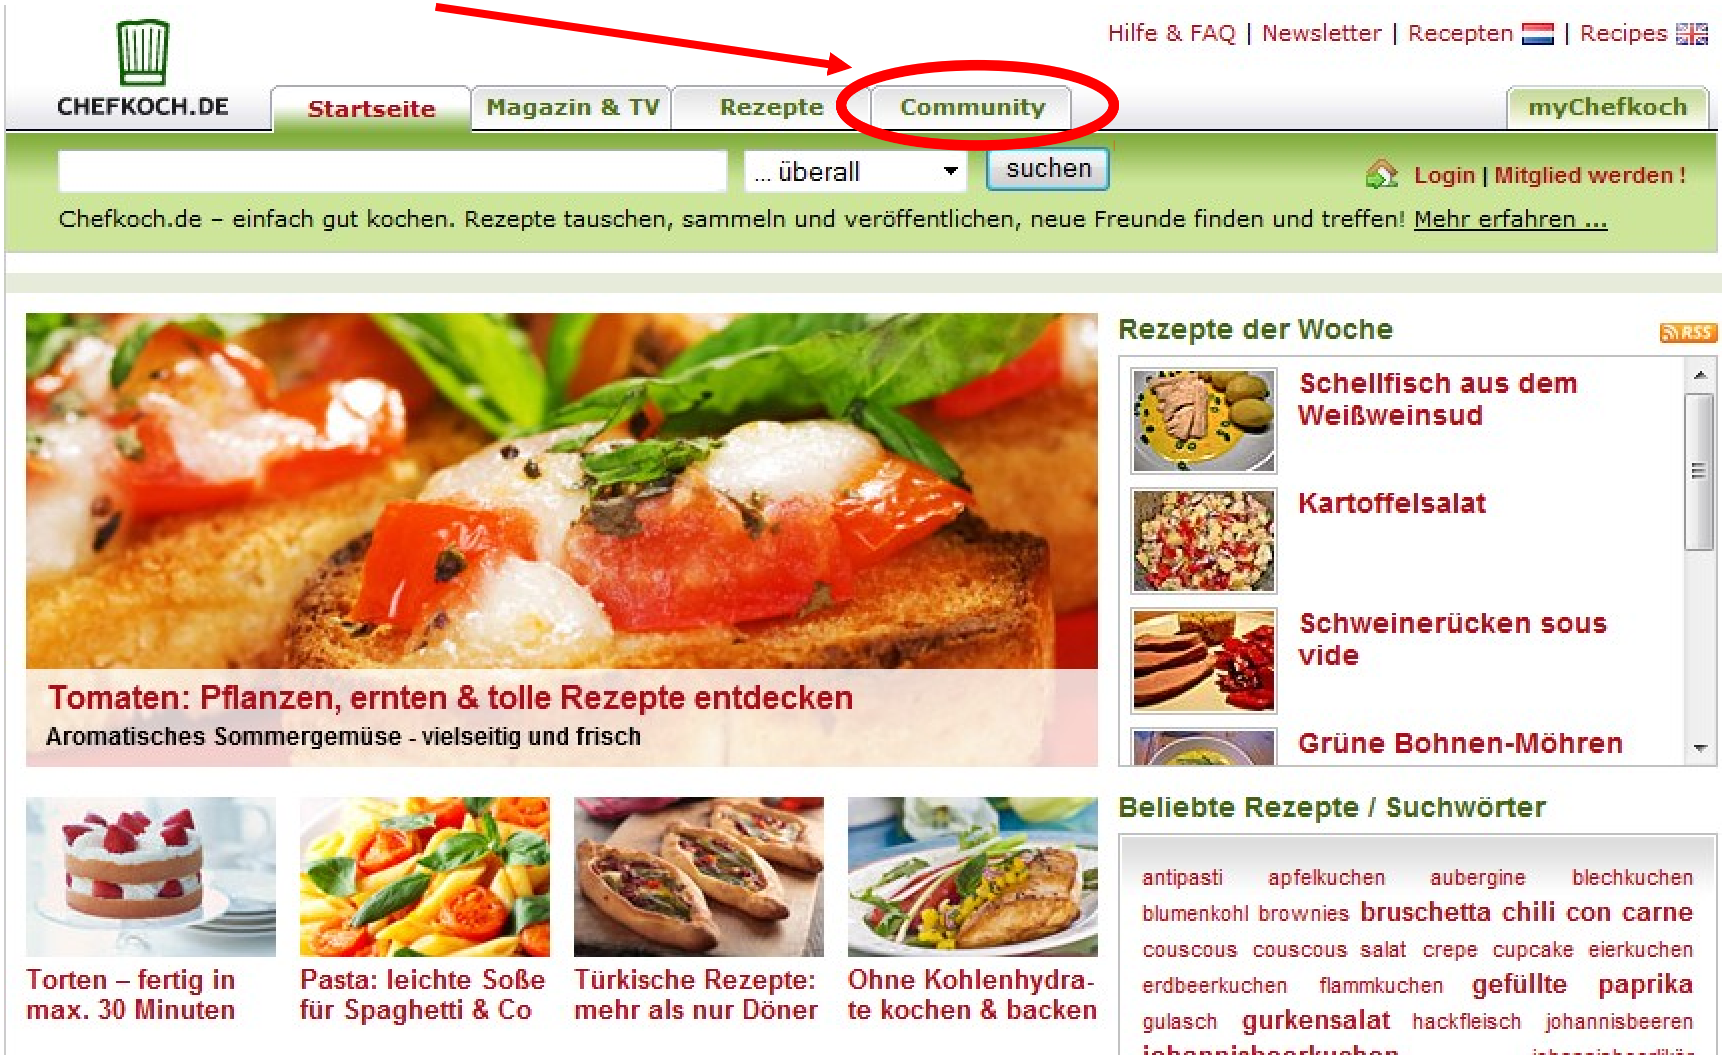

Domain-Specific Corpus

Runde Zucchini

Vom 07.07.2012 21:43

Fietelino

Mitglied seit 29.10.2008
84 Beiträge (ø0,06/Tag)

Hallo liebe Chefköche,

bin auf der Suche nach Rezepten mit runden Zucchini. Habe hier in der Datenbank nur ein Rezept gefunden. Auch würde mich interessieren wo man die runden Zucchini zu kaufen bekommt. Habe sie noch nicht wirklich gesehen. Habe nur im Vorbeigehen in Zeitschriften gesehen, das es sie gibt.

Vielen Dank im Vorraus für Eure Hilfe.

LG Birgit

Diesen Beitrag / einen Verstoß melden

Vom 07.07.2012 21:49

Mojoverde

Mitglied seit 14.11.2008
7.947 Beiträge (ø5,95/Tag)

hallo
wo Du sie in Hamburg bekommst, weiss ich nicht.
Sie eignen sich zum Füllen mit Hackfleisch, evtl käsehaltig und zum Überbacken.
In Scheiben geschnitten, paniert und ausgebacken mit Tomatensauce schmecken sie auch sehr gut.
gruß mo

Diesen Beitrag / einen Verstoß melden

Vom 07.07.2012 21:53

noch kein Foto

kirsten2371

Hallo Birgit,

runde Zucchinis kannst Du doch eigentlich wie normale Zucchinis verwenden nur die Garzeiten würde ich eventuell je nach dem bis 10 Minuten verlängern.

Distant Supervision Labeling

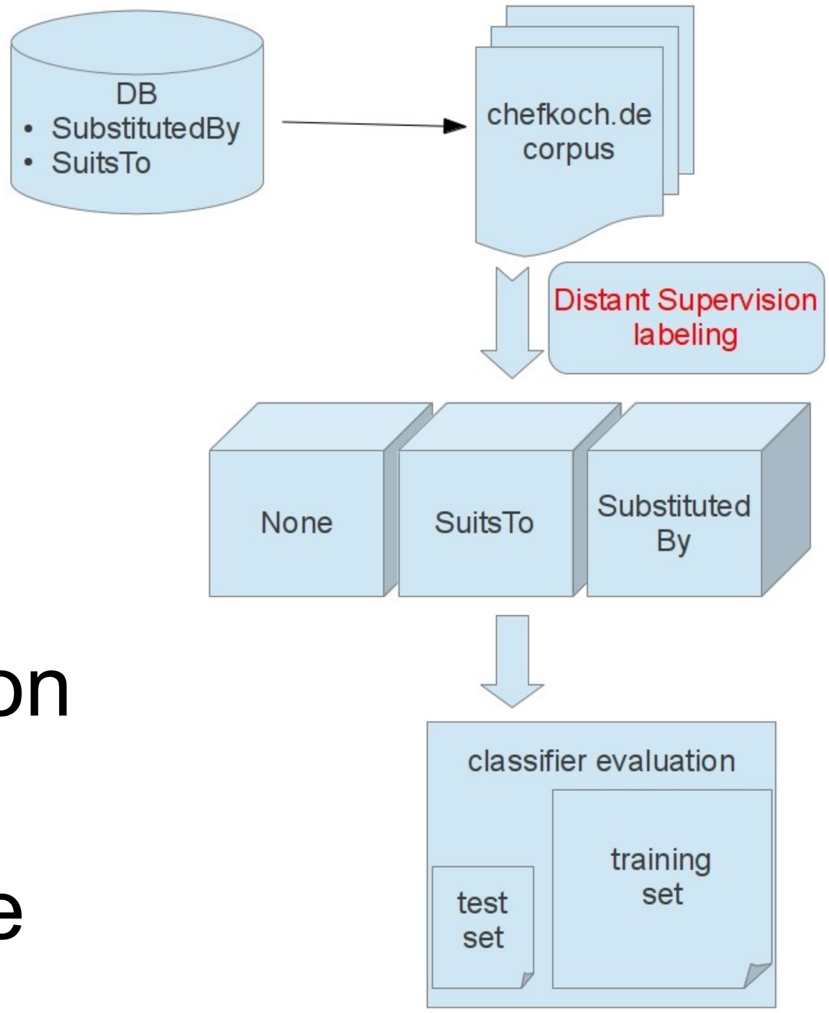

Using:

- Distant Supervision Assumption
- Relation database

Distant Supervision Labeling

DB of argument pairs,
e.g. <hamburger, fries>

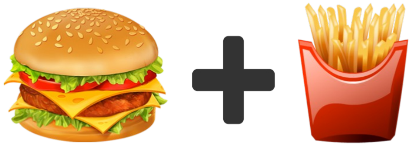

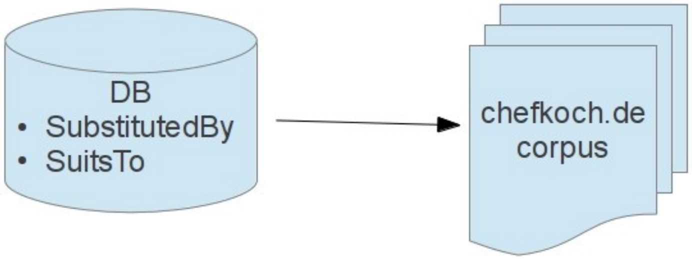

Distant Supervision Labeling

DB of argument pairs,
e.g. <hamburger, fries>

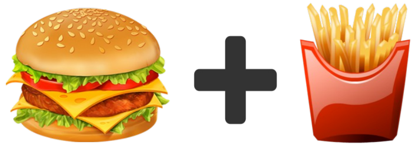

Look for sentences:
“My children's favourite is **hamburger** with **fries**.“
→ label: SuitsTo

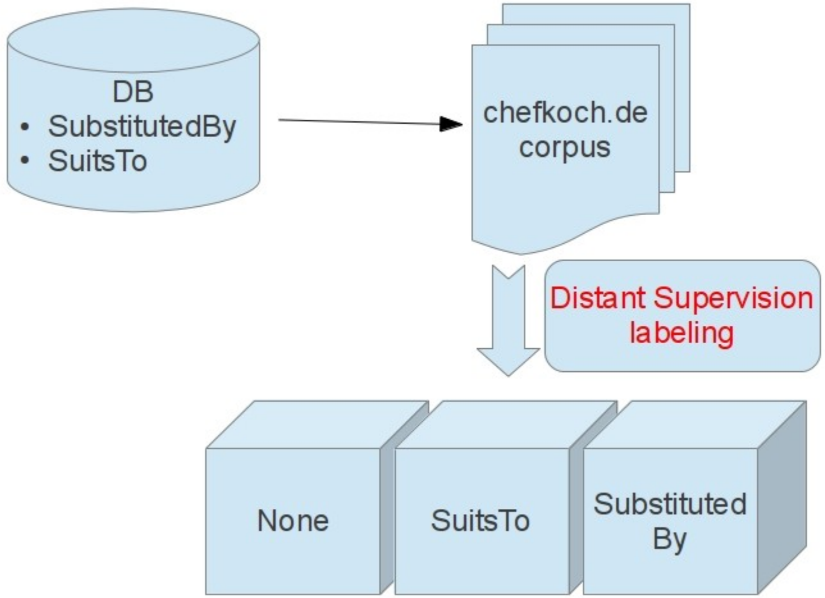

Distant Supervision Labeling

DB of argument pairs,
e.g. <hamburger, fries>

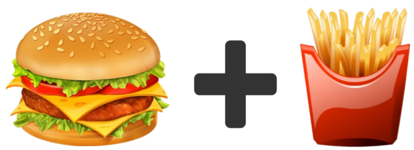

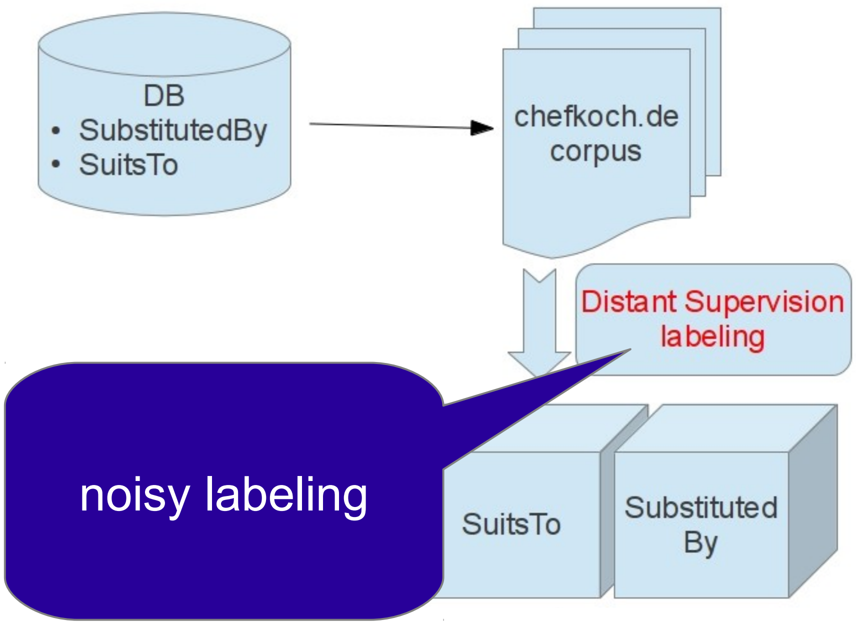

Look for sentences:
“My children's favourite is **hamburger** with **fries**.“
→ label: SuitsTo

Distant Supervision Labeling

DB of argument pairs,
e.g. <hamburger, fries>

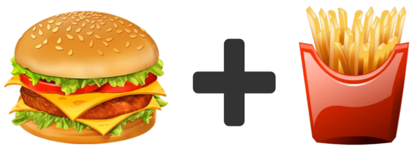

Look for sentences:
“My children's favourite is **hamburger** with **fries**.“
→ label: SuitsTo

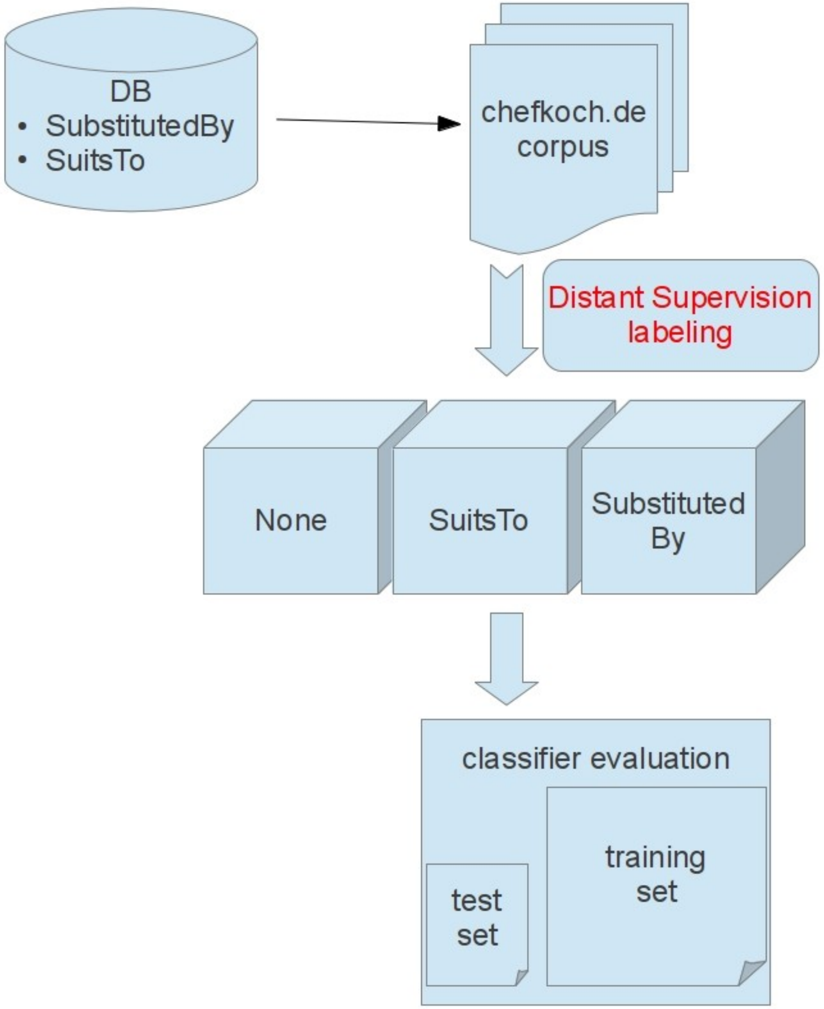

Distant Supervision Labeling

Gold Standard Testset:

- 2240 sentences, randomly chosen from corpus.
- Manually labeled with food relations (SuitsTo, SubstitutedBy, or None).

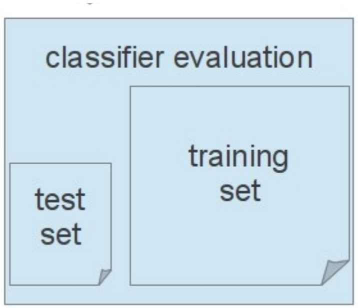

Distribution of Classes

- ***In the gold standard:***

Relation	Example	class ratio in %
SuitsTo	My kids love the simple combination of fish fingers with mashed potatoes.	60
SubstitutedBy	We usually buy margarine instead of butter.	9
None	On my shopping list, I've got bread, cauliflower, ...	31

- ***In the unlabeled food corpus:***

Relation	Argument pairs matched	Sentences matched
SuitsTo	1,374	44,692
SubstitutedBy	789	34,771
None	62,191	1,187,101

Outline of Talk

- Introduction
- Tasks & Data
- Methods and Experiments of Producing Training Data
- Conclusions

Setup for Relation Extraction

- 10K training sentences in each configuration.
- Train on distantly-labeled data.
- Enforce estimated class distribution.
- Train Support Vector Machines.
- Binary classifier for each of the target relation types (i.e., SuitsTo, SubstitutedBy).
- Starting with standard feature set (bag-of-words, POS, n-grams, etc.)

Degrees of Freedom in Building Classifiers

Relation Types	**Knowledge**	**Processing Level**	**Classifier**
SuitsTo SubstitutedBy	patterns types	argument level sentence level feature level	DS rule-based

Degrees of Freedom in Building Classifiers

Relation Types	Knowledge	Processing Level	Classifier
SuitsTo SubstitutedBy	patterns types	argument level sentence level feature level	DS rule-based

Knowledge: Patterns

<foodItem_1> . . . <foodItem_2>

Relation	Pattern	Example
SuitsTo	FOOD **and** FOOD, FOOD **with** FOOD, FOOD **fit to** FOOD	My kids love the simple combination of fish fingers **with** mashed potatoes.
SubstitutedBy	FOOD **or** FOOD, FOOD **instead of** FOOD, ...	We usually buy margarine **instead of** butter.

Knowledge: Food Types

Food Guide Pyramid (U.S. Department of Agriculture, 1992)

Knowledge: Food Types

Food Guide Pyramid (U.S. Department of Agriculture, 1992)

- Groups food items into 11 common categories.
- Items within the same category share similar (nutritional) properties.
- semi-automatic mapping

(Wiegand et al. 2014)

Knowledge: Food Types

Food Guide Pyramid (U.S. Department of Agriculture, 1992)

example category:

fruits

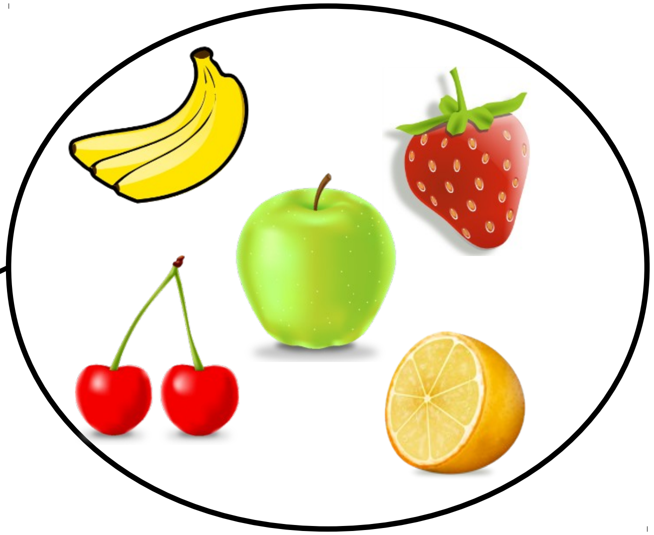

Knowledge: Food Types

Food Guide Pyramid (U.S. Department of Agriculture, 1992)

example category:

meat & fish

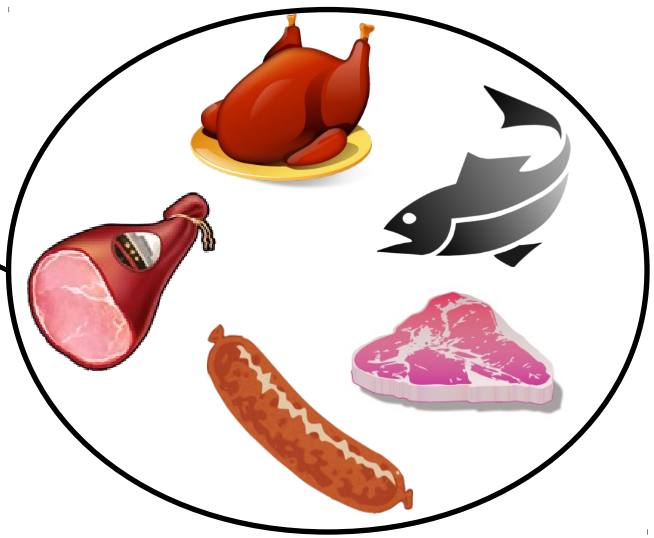

Pattern Ambiguity

Pattern <*foodItem_1*> ***and*** <*foodItem_2*>:

- I very often eat fish and fries.

→ *SuitsTo*

- For these types of dishes you can offer both, Burgundy wine and Champagne.

→ *SubstitutedBy*

Pattern Ambiguity

Pattern *<foodItem_1> and <foodItem_2>*:

- I very often eat fish[MEAT] and fries[STARCH].

→ *SuitsTo*

- For these types of dishes you can offer both, Burgundy wine[BEVER] and Champagne[BEVER].
  → *SubstitutedBy*

How can food categories help resolve this ambiguity?

Pattern Ambiguity

Pattern <*foodItem_1*> *and* <*foodItem_2*>:

- I very often eat fish[MEAT] and fries[STARCH].

→ *SuitsTo*

- For these types of dishes you can offer both Burgundy wine[BEVER] and Champagne[BEVER].
  → *SubstitutedBy*

Type Assumption:

Relation	Rule
SuitsTo	<MEAT, STARCH>
SubstitutedBy	<BEVERAGE, BEVERAGE>

Pattern Ambiguity

Pattern *<foodItem_1> and <foodItem_2>*:

- I very often eat fish[MEAT] and fries[STARCH].

→ *SuitsTo*

- For these types of dishes you can offer both Burgundy wine[BEVER] and Champagne[BEVER].
  → *SubstitutedBy*

Type Assumption:

Relation	Rule
SuitsTo	<**X, Y**>
SubstitutedBy	<**X, X**>

Degrees of Freedom in Building Classifiers

Relation Types

SuitsTo
SubstitutedBy

Knowledge

patterns
types

Processing Level

argument level
sentence level
feature level

Classifier

DS
rule-based

Argument-Level Data Selection

Idea:
Use **typical** positive (negative)
relation instances for training.

→ Create ranking of argument pairs.
- High ranks: positive instances.
- Low ranks: negative instances.

→ Select argument pairs from these rankings.

Argument-Level Data Selection

Except from a random selection of argument pairs, consider:

external knowledge	statistical association
patterns	frequency
food types	PMI
ontology-based similarity (WordNet)	

→ Create rankings.
→ Select argument pairs from these rankings.

Argument-Level Data Selection: Results

Using the standard feature set:

	SuitsTo						SubstitutedBy					
	positive						*positive*					
negative	rand.	freq	pmi	ontol.	patt	type	rand.	freq	pmi	ontol.	patt	type
random												
freq												
pmi												
ontol.												
patt												
type												

Argument-Level Data Selection: Results

Using the standard feature set:

	SuitsTo						SubstitutedBy					
	positive						*positive*					
negative	rand.	freq	pmi	ontol.	patt	type	rand.	freq	pmi	ontol.	patt	type
random												
freq												
pmi												
ontol.												
patt												
type												

Conditions:
- Distant Supervision labeling
- SVM binary classifiers
- Standard feature set
- Gold standard evaluation

Results reported in F-score

Argument-Level Data Selection: Results

Using the standard feature set:

	SuitsTo						SubstitutedBy					
	positive						*positive*					
negative	rand.	freq	pmi	ontol.	patt	type	rand.	freq	pmi	ontol.	patt	type
random	41.8						61.8					
freq												
pmi												
ontol.												
patt			**60.4**								**66.8**	
type												

→ Overall, data selection *does* make a difference.

Argument-Level Data Selection: Results

Best methods for choosing positive instances:

	SuitsTo						SubstitutedBy					
	positive						*positive*					
negative	rand.	freq	pmi	ontol.	patt	type	rand.	freq	pmi	ontol.	patt	type
random	41.8		49.2				61.8				65.1	
freq			50.1								64.0	
pmi			50.7								64.8	
ontol.			50.3								64.7	
patt			**60.4**								**66.8**	
type			49.2								64.6	

→ Effectiveness of methods differs between relation types.

Argument-Level Data Selection: Results

	SuitsTo						SubstitutedBy					
	positive						*positive*					
negative	rand.	freq	pmi	ontol.	patt	type	rand.	freq	pmi	ontol.	patt	type
random	41.8		49.2				61.8				65.1	
freq			50.1								64.0	
pmi			50.7								64.8	
ontol.			50.3								64.7	
patt			60.4								66.8	
type			49.2								64.6	

→ SuitsTo: PMI as measure of argument association.

→ SubstitutedBy: high-precision patterns.

Sentence-Level Data Selection

Idea:
Use relation-specific knowledge to filter training instances:

external knowledge:

patterns

food types

Select sentences that:
→ Match any pattern of the target relation.
→ Satisfy the type assumption.

Sentence-Level Data Selection: Results

Using the standard feature set, and the best configuration from argument-level filtering:

SuitsTo				SubstitutedBy			
	positive				*positive*		
negative	**no filter**	**pattern**	**type**	*negative*	**no filter**	**pattern**	**type**
no filter				**no filter**			
pattern				**pattern**			
type				**type**			

Sentence-Level Data Selection: Results

Using the standard feature set, and the best configuration from argument-level filtering:

SuitsTo				SubstitutedBy			
	positive				*positive*		
negative	**no filter**	**pattern**	**t**	*negative*	**no**	**pattern**	**type**
no filter							
pattern							
type							

- Distant Supervision labeling
- SVM binary classifiers
- Standard feature set
- *Pre-Selection: best config. from argument-level filtering*
- Gold standard evaluation

Results reported in F-score

Sentence-Level Data Selection: Results

Using the standard feature set, and the best configuration from argument-level filtering:

SuitsTo				SubstitutedBy			
	positive				*positive*		
negative	**no filter**	**pattern**	**type**	*negative*	**no filter**	**pattern**	**type**
no filter	60.40			**no filter**	66.80		
pattern				**pattern**			
type				**type**			

Sentence-Level Data Selection: Results

Using the standard feature set, and the best configuration from argument-level filtering:

SuitsTo				SubstitutedBy			
	positive				*positive*		
negative	**no filter**	**pattern**	**type**	*negative*	**no filter**	**pattern**	**type**
no filter	60.40	43.22	60.66	**no filter**	66.80	64.52	66.05
pattern	47.44	49.19	47.51	**pattern**	28.34	63.66	27.05
type	60.70	43.15	60.90	**type**	67.53	64.50	66.15

→ Sentence-level filtering does not help.

Feature-Level Processing

Idea:
Use relation-specific knowledge as features.

external knowledge:

patterns

food types

Additional features encode:
→ Whether a sentence matches a pattern.
→ Type of relation arguments.

Feature-Level Processing: Results

Adding pattern and type information to the best configuration from previous filtering:

	argument level		sentence level	feature level		
	random	best	best	+pattern	+type	all
SuitsTo	41.78	60.40	60.90			
SubstitutedBy	61.75	66.80	67.53			

Feature-Level Processing: Results

Adding pattern and type information to the best configuration from previous filtering:

	argument level	sentence level	feature level	
				all
SuitsTo				
SubstitutedB				

- Distant Supervision labeling
- SVM binary classifiers
- Standard feature set
- *Pre-Selection: best config. from argument- & sentence-level filtering*
- Gold standard evaluation

Results reported in F-score

Feature-Level Processing: Results

Adding pattern and type information to the best configuration from previous filtering:

	argument level		sentence level	feature level		
	random	best	best	+pattern	+type	all
SuitsTo	41.78	60.40	60.90	60.58	61.81*	**61.89***
SubstitutedBy	61.75	67.00	67.53	67.78	70.37*	**70.50***

* significantly better than best sentence-level result at $p < 0.05$ (paired t-test)

Feature-Level Processing

Adding pattern and type information to the best configuration from previous filtering:

	argument level		sentence level	feature level		
	random	best	best	+pattern	+type	all
SuitsTo	41.78	60.40	60.90	60.58	61.81*	**61.89***
SubstitutedBy	61.75	67.00	67.53	67.78	70.37*	**70.50***

* significantly better than best sentence-level result at $p < 0.05$ (paired t-test)

→ Type information is beneficial.

Feature-Level Processing

Adding pattern and type information to the best configuration from previous filtering:

	argument level		sentence level	feature level		
	random	best	best	+pattern	+type	all
SuitsTo	41.78	60.40	60.90	60.58	61.81*	**61.89***
SubstitutedBy	61.75	67.00	67.53	67.78	70.37*	**70.50***

* significantly better than best sentence-level result at $p < 0.05$ (paired t-test)

→ Type information is beneficial.

→ Pattern information is not.

Feature-Level Processing

Adding pattern and type information to the best configuration from previous filtering:

	argument level		sentence level	feature level		
	random	best	best	+pattern	+type	all
SuitsTo	41.78	60.40	60.90	60.58	61.81*	**61.89***
SubstitutedBy	61.75	67.00	67.53	[illegible]	[illegible]	**70.50***

Remember: patterns **did** help on argument-level filtering.

→ Type information is beneficial.
→ Pattern information is not.
→ Impact of knowledge depends on processing level.

Degrees of Freedom in Building Classifiers

Relation Types

SuitsTo
SubstitutedBy

Knowledge

patterns
types

Processing Level

argument level
sentence level
feature level

Classifier

DS
rule-based

Classifier Comparison

- Distantly-supervised learner
- Simplistic, rule-based decision

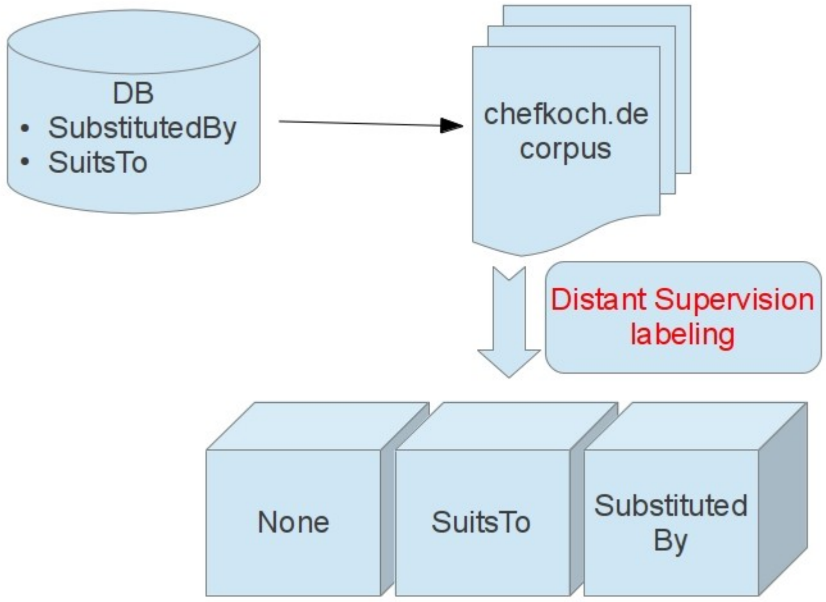

Classifier Comparison

- Distantly-supervised learner
- Simplistic, rule-based decision

Relation	Pattern-Rule
SuitsTo	**X and Y, X with Y, ...**
SubstitutedBy	**X or Y, X instead of Y, ...**

Relation	Type-Assumption
SuitsTo	<**X, Y**>
SubstitutedBy	<**X, X**>

Feature-Level Processing

Comparing classifiers:
Distant Supervision (DS) vs. rule-based (RB)

	Major	DS_{random}	DS_{best}	$RB_{pattern}$	RB_{type}	super
SuitsTo	37.50	41.78	61.89	42.60	62.65	73.48
SubstitutedBy	47.64	61.75	70.51	63.97	62.02	77.77

Feature-Level Processing

Comparing classifiers:
Distant Supervision (DS) vs. rule-based (RB)

	Major	DS_{random}	DS_{best}	$RB_{pattern}$	RB_{type}	super
SuitsTo	37.50	41.78	61.89	42.60	62.65	73.48
SubstitutedBy	47.64	61.75	70.51	63.97	62.02	77.77

→ SuitsTo: RB achieves competitive (better) results.
→ Type information seems highly informative.

Feature-Level Processing

Comparing classifiers:
Distant Supervision (DS) vs. rule-based (RB)

	Major	DS_{random}	DS_{best}	$RB_{pattern}$	RB_{type}	super
SuitsTo	37.50	41.78	61.89	42.60	62.65	73.48
SubstitutedBy	47.64	61.75	70.51	63.97	62.02	77.77

→ SubstitutedBy: DS learner outperforms RB.
→ DS learner profits from multiple information sources.

Outline of Talk

- Introduction
- Tasks & Data
- Methods and Experiments of Producing Training Data
- Experiments
- Conclusions

Conclusions (I)

Findings regarding potential improvement of Distant Supervision learning:

- Argument-level and feature-level processing help to increase classification performance.
- Sentence-level filtering is not beneficial.

Conclusions (I)

Findings regarding potential improvement of Distant Supervision learning:

- Argument-level and feature-level processing help to increase classification performance.
- Sentence-level filtering is not beneficial.
- Effectiveness of external knowledge varies, depending on
  - the relation type,
  - the processing level.
- Patterns beneficial on argument-level.
- Food types beneficial on feature-level.

Conclusions (II)

- Substantial improvement of performance by
  - careful selection of training data,
  - appropriate feature design.

Conclusions (II)

- Substantial improvement of performance by
  - careful selection of training data,
  - appropriate feature design.

- Simple rule-based decision achieves competitive performance in some cases.

Thank You!

References (I)

- ***The Distant Supervision Assumption:***

  Mintz, M., Bills, S., Snow, R., Jurafsky, D.: Distant Supervision for Relation Extraction without Labeled Data. In: Proc. of ACL/IJCNLP (2009).
- **The *chefkoch.de* corpus, the surface patterns:**

  Wiegand, M., Roth, B., Klakow, D.: Web-based Relation Extraction for the Food Domain. In: Bouma, G., Ittoo, A., Métais, E., Wortmann, H. (eds.) NLDB 2012. LNCS, vol. 7337, pp.222-227. Springer, Heidelberg (2012a).
- ***The relation database, annotation schema:***

  Wiegand, M., Roth, B., Lasarcyk, E., Köser, S., Klakow, D.: A Gold Standard for Relation Extraction in the Food Domain. In: Proc. of LREC (2012b).

References (II)

- ***The Food Guide Pyramid:***

  U.S. Department of Agriculture, H.N.I.S.: *Home and Garden Bulletin, 252*, Washington D.C., USA (1992).
- ***Food Types:***

  Wiegand, M., Roth, B., Klakow, D.: Automatic Food Categorization from Large Unlabeled Corpora and Its Impact on Relation Extraction. In: Proc. of EACL (2014).
- ***Ontological Similarity Measure:***

  Wu, Z., Palmer, M.: Verbs semantics and lexical selection. In: Proc. of ACL (1994).

  Hamp, B., Feldweg, H.: GermaNet – a Lexical-Semantic Net for German. In: Proc. of ACL Workshop. Automatic Information Extraction and Building of Lexical Semantic Resources for NLP Applications (1997).

BACK-UP SLIDES

Standard Feature Set

Feature	Description
word-left-window	a window of 2 words to the left of *arg1*
word-right-window	a window of 2 words to the right of *arg2*
word-window	the word sequence between *arg1* and *arg2*
left-lemma-window	a window of 2 words to the left of *arg1* as lemmas
right-lemma-window	a window of 2 words to the right of *arg2* as lemmas
lemma-window	the word sequence between *arg1* and *arg2* as lemmas
bow	all words in the sentence
lemma-bow	lemmas of w_{i-1}, w_{i+1} where $w_i \in \{arg1, arg2\}$
lemma-bigrams	all bigrams $< w_i, w_j >$ between *arg1* and *arg2* as lemmas
pos-left-window	part-of-speech tags of words in a window of 2 words to the left of *arg1*
pos-right-window	part-of-speech tags of words in a window of 2 words to the right of *arg2*
pos-window	part-of-speech sequence between *arg1* and *arg2*
pos-unigrams	part-of-speech tags of w_{i-1}, w_{i+1} where $w_i \in \{arg1, arg2\}$
pos-bigrams	all part-of-speech bigrams $< t_i, t_j >$ between *arg1* and *arg2* using lemmas

Gold Standard

- 2240 manually-labeled sentences.
- “Natural” class distribution estimated from a labeled random sample of 100 sentences.
- Annotation schema from (Wiegand et al., 2012b).
- Inter-annotator agreement

  - on a subset of 400 sentences

  - 2 annotators

  - Kappa 0.78 (Cohen’s Kappa (Cohen, 1960)

  → substantial agreement (Landis and Koch,1977).

Data Selection Methods

Level	Method	Description
argument	random	select argument pairs from relation database at random
	frequency	sort argument pairs according to frequency (of co-occurring in the text corpus)
	pmi	sort argument pairs according to *pointwise mutual information (pmi)*
	patt^{+}	sort argument pairs according to pmi of food items and the surface patterns pertaining to the target relation in *descending* order
	patt^{-}(a)	sort argument pairs according to pmi of food items and the surface patterns pertaining to the target relation in *ascending* order
	patt^{-}(b)	sort argument pairs according to pmi of food items and the surface patterns pertaining to the contrast relation in *descending* order
	type	sort type pairs (e.g. <MEAT,STARCH>) according to pmi and consider their actual food instantiations as arguments
	wup	sort argument pairs according to Wu-&-Palmer [8] similarity in GermaNet [9]
sentence	pattern	only include sentences in which target food items co-occur with surface pattern from target relation
	type	only include sentences in which type rule for the pertaining relation (Table 5) is fulfilled
feature	pattern	include all surface patterns as additional features
	type	include features indicating the types of the target food items, e.g. <MEAT,STARCH> for <*fish*,*chips*>
	standard	standard features directly extracted from training data without external knowledge resources (see Table 7)

Argument-Level Data Selection

Using the standard feature set:

	SuitsTo						SubstitutedBy					
	positive						*positive*					
negative	**rand.**	**freq**	**pmi**	**wup**	**patt**$^+$	**type**	**rand.**	**freq**	**pmi**	**wup**	**patt**$^+$	**type**
random	41.8	45.0	49.2	46.0	45.3	42.1	61.8	59.0	63.0	59.8	65.1	60.0
freq	40.1	44.1	50.1	41.4	44.8	40.4	61.0	58.1	61.7	59.0	64.0	59.0
pmi	42.3	45.1	50.7	43.5	47.8	43.0	62.8	58.9	64.2	61.0	64.8	59.3
wup	42.4	45.8	50.3	44.4	45.8	42.0	59.3	57.4	60.8	57.6	64.7	55.3
patt$^-$ **(a)**	43.7	47.2	52.2	45.4	47.7	42.0	64.8	62.6	**67.0**	62.9	66.8	63.5
patt$^-$ **(b)**	55.0	56.6	**60.4**	56.7	54.9	54.5	52.5	53.1	57.8	54.7	62.2	50.3
type	42.4	43.8	49.2	44.3	45.9	41.0	61.5	59.3	63.3	59.0	64.6	61.1

→ SuitsTo: PMI as measure of argument association.

→ SubstitutedBy: high-precision patterns.

Feature-Level Processing

Comparing classifiers:
(training on best feature configuration)

		Major	$\mathbf{DS}_{random}$	$\mathbf{DS}_{best}$	$\mathbf{RB}_{pattern}$	$\mathbf{RB}_{type}$	$\mathbf{RB}_{comb}$	**super**
SuitsTo	**Acc**	60.00	60.84	**68.05**	45.49	68.03	44.73	77.58
	Prec	30.00	53.74	**67.45**	55.40	66.91	59.03	81.32
	Rec	50.00	50.59	62.18	53.13	**62.69**	53.66	72.49
	F	37.50	41.78	61.89	42.60	**62.65**	40.00	73.48
SubstitutedBy	**Acc**	91.00	86.39	86.86	91.47	79.28	**91.56**	92.54
	Prec	45.50	61.28	67.29	77.72	60.55	**81.72**	78.87
	Rec	50.00	62.33	**77.59**	60.28	72.13	57.99	77.31
	F	47.64	61.75	**70.51**	63.97	62.02	61.26	77.77

→ SuitsTo: type information seems highly informative.
→ SubstitutedBy: profits from multiple information sources.

Why is RB sometimes better than RB?

In general:

- RB classifiers are high-precision classifiers.
- Learners are high-recall classifiers.

- SuitsTo is a majority class → a good classifier should aim for high precision.
- SubstitutedBy is a minority class → a good classifier should aim for high recall.

- → For SuitsTo, the RB classifier is a suitable match.