Stochastic Models for Semantic Parsing, Multi-Faceted Topic Discovery, and Causal Event Inference: Perspectives from Natural Language Processing

ROXANA GIRJU
Linguistics and Computer Science,
Beckman Institute,
University of Illinois

Semantic Frontiers' Group: Rania Al-Sabbagh, Brendon Beamer, Chen Li, David Lundgren, Michael Paul (now at JHU), Mehwish Riaz

November 12, 2011

Some Research Projects @ *Semantic Frontiers* Group

- P1. Stochastic Models for Semantic Parsing
  - The problem of knowledge discovery
  - Semantic relation discovery
  - Some text-image applications
- P2. Perspectives, aspects and sentiment
  - in scientific literature
  - in Israeli-Palestinian editorials
  - cultural differences from travelers' experiences
- P3. Causal Event Inference
  - Identification of causal relations between events
  - Applications: question answering, textual entailment

Semantic Parsing (1)

- Knowledge Discovery from Text:
  - The process of extracting useful, non-trivial (implicit) knowledge from unstructured data.

- Knowledge Discovery as Semantic Relations:
  - are underlying relations between two concepts expressed by words or phrases
  - Examples:
    - HYPERNYMY (IS-A),
    - MERONYMY (PART-WHOLE),
    - CAUSE - EFFECT, etc.

- Semantic parsing:
    - supports automated reasoning.

Semantic Parsing (2)

Unstructured Information

Web
Documents
News
Digital library

Semantic Parser

Structured Knowledge

- KB
- Semantically tagged text
  - Concepts
  - Semantic relations
  - Links between multiple docs

The task of semantic relation discovery:

Given a pair of nouns n1 – n2, determine the pair's meaning.

Semantic Parsing (3)

Example 1: (Girju et al. 2003, 2006, 2007, 2009)

[*Saturday's snowfall*]TEMP topped [*a record* in *Hartford, Connecticut*]LOC with [*the total of 12.5 inches*]MEASURE, [the *weather service*]TOPIC said. The storm claimed its fatality Thursday when [a *car* driven by a [*college student*]PART-WHOLE]THEME skidded on [*an interstate overpass*]LOC in [*the mountains of Virginia*]LOC/PART-WHOLE and hit [*a concrete barrier*]PART-WHOLE, police said.

(www.cnn.com – "Record-setting Northeast snowstorm winding down", December 7, 2003)

TEMP (Saturday, snowfall)
LOC (Hartford Connecticut, record)
MEASURE (total, 12.5 inch)
TOPIC (weather, service)
PART-WHOLE (student, college)
THEME (car, driven by a college student)

LOC (interstate, overpass)
LOC (mountains, Virginia)
PART-WHOLE/LOC (mountains, Virginia)
PART-WHOLE (concrete, barrier)

Semantic Parsing: Basic Approach (1)

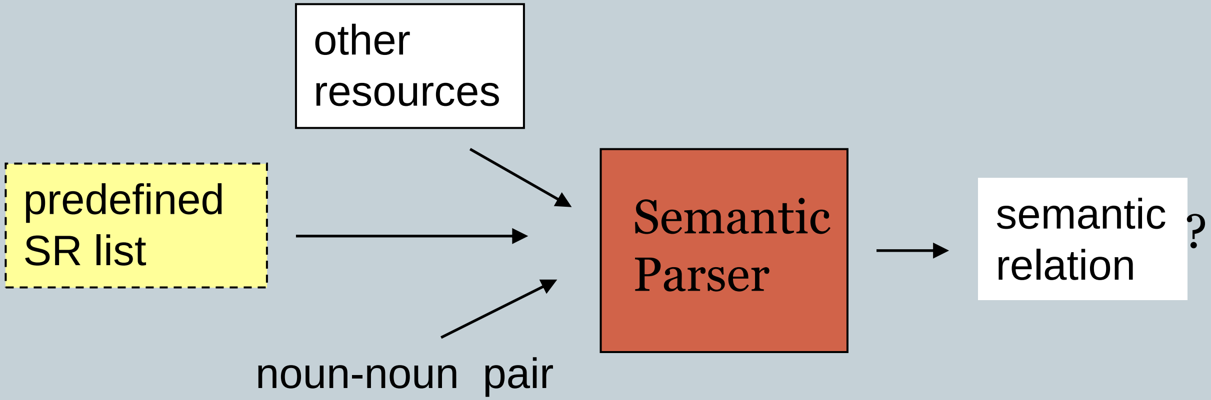

Semantic Parsing: Basic Approach (2)

Examples of relations (SemEval 2007):

- **Cause-Effect**: laugh wrinkles
- **Instrument-Agency**: laser printer
- **Product-Producer**: honey bee
- **Origin-Entity**: message from outer-space
- **Theme-Tool**: news conference
- **Part-Whole**: car door
- **Content-Container**: the cookies in the jar

Semantic Parsing: Basic Approach (3)

- (SemEval 2007):

After the cashier put the <e1>**cash**</e1> in a <e2>**bag**</e2>, the robber saw a bottle of scotch that he wanted behind the counter on the shelf.

Query = “the * in a *”
WordNet(e1) = "cash%1:21:00::"
WordNet(e2) = "bag%1:06:00::"
Content-Container(e1,e2) = “true”

A Stochastic Model for Semantic Parsing

Semantic Scattering2 (SS2)

(Moldovan at al. 2004) (Badulescu & Moldovan 2005) (Girju at al. 2005) (Beamer, Girju, Rozovskaya, 2008)(Girju et al 2010)

- The most important component of our full Semantic Parser @SemEval 2007
- Input: n1 (and its sense in context);
n2 (and its sense in context);
list of semantic relations;
WordNet noun hierarchy
- Output: <n1, n2, r>

A Stochastic Model for Semantic Parsing

Datasets

- 140 training and >70 test examples for each relation;
- Balanced positive and negative examples.
- Definition provided for each relation
- WordNet senses provided for input nouns

- Annotation summary:
  - high inter-annotator agreement on WordNet senses and semantic relations
  - Disagreements discussed and consensus reached (or example thrown out).

A Stochastic Model for Semantic Parsing

Hypothesis:

- Noun – noun pairs with the same/similar meaning tend to encode the same semantic relation.

Approach:

- The semantic class of a noun:
  - specifies its WordNet sense in context and
  - implicitly points to its hypernyms;

A Stochastic Model for Semantic Parsing

```
Sense 3
bass, basso --
(an adult male singer with the lowest voice)
=> singer, vocalist
   => musician, instrumentalist, player
      => performer, performing artist
         => entertainer
            => person, individual, someone...
               => life form, organism, being...
                  => entity, something
               => causal agent, cause, causal agency
                  => entity, something

Sense 7
bass --
(the member with the lowest range of a family of
musical instruments)
=> musical instrument
   => instrument
      => device
         => instrumentality, instrumentation
            => artifact, artefact
               => object, physical object
                  => entity, something
```

A Stochastic Model for Semantic Parsing

Features (for binary classifiers):

- Semantic class of <u>head noun</u>: f_j^h
- Semantic class of <u>modifier noun</u>: f_i^m
  - *E.g.: hand#1 of a woman#1* [P-W]
- Feature pair: $<f_i^m, f_j^h>= f_{ij}$
- Form tuples: $<f_{ij}, r>$

$$P(r \mid f_{ij}) = \frac{n(r, f_{ij})}{n(f_{ij})}$$

$$\hat{r} = \arg\max P(r \mid f_{ij})$$

A Stochastic Model for Semantic Parsing

The task: Find the best set of semantic classes (i.e., a boundary G* – a division in WordNet) that best generalize over the training data and accurately classify unseen Data.

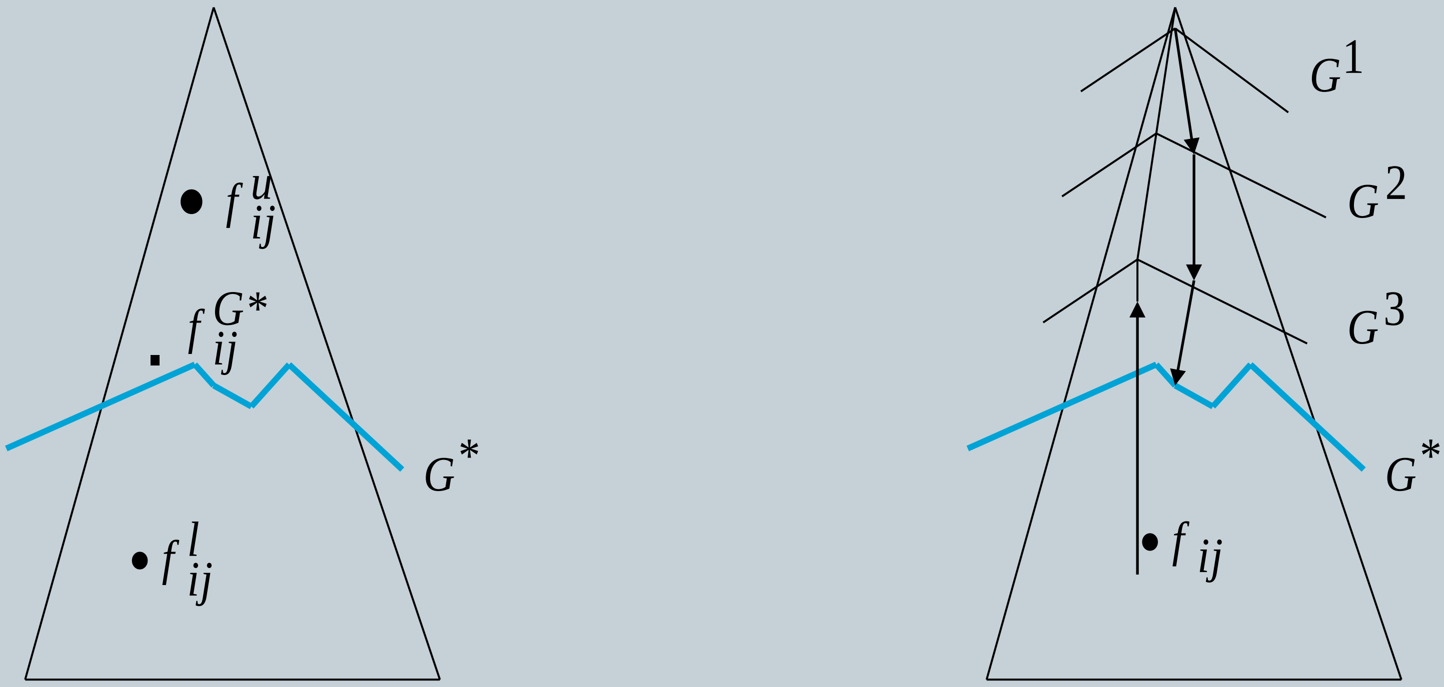

A Stochastic Model for Semantic Parsing

Example for Part-Whole

- Step 1: (Create an Initial Boundary (i.e., generalize the training examples):
  - Initial corpus:

    <n1#sense; n2#sense; target>

    E.g.: <hand#1; woman#1; YES>

<hand#1, **entity**#1; woman#1, **entity**#1; YES>

A Stochastic Model for Semantic Parsing

- Ambiguous examples:

<n1_hierarchy#sense; n2_hierarchy#sense; Yes/No>

<apartment #1; woman #1; No>
<hand #1; woman #1; Yes>

<**entity** #1; **entity** #1; Yes/No>

A Stochastic Model for Semantic Parsing

Step 2: Specialize ambiguous examples:

woman#1 ‘s apartment#1;
hand#1 of a woman#1;

<entity#1; entity#1; Yes/No>

specialization

<whole#2; causal_agent#1; No>
<part#7; causal_agent#1; Yes>

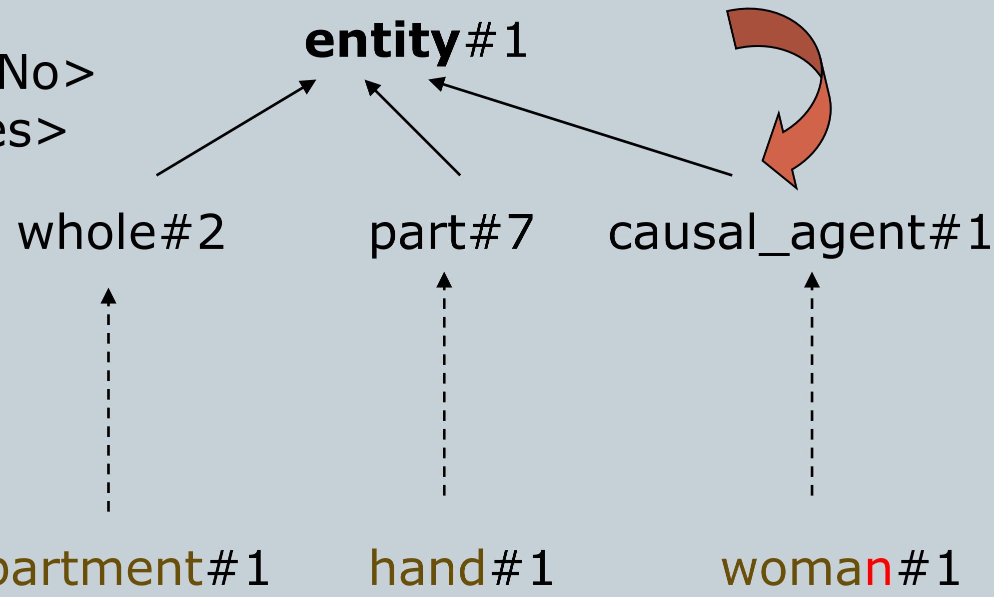

A Stochastic Model for Semantic Parsing

Specialization example2:

<entity#1; entity#1; Yes/No>

specialization

<part#7; organism#1; Yes>
<part#7; organism#1; No>

specialization

. . .

leg#2 of insect#1;
insect#1 's world#7;

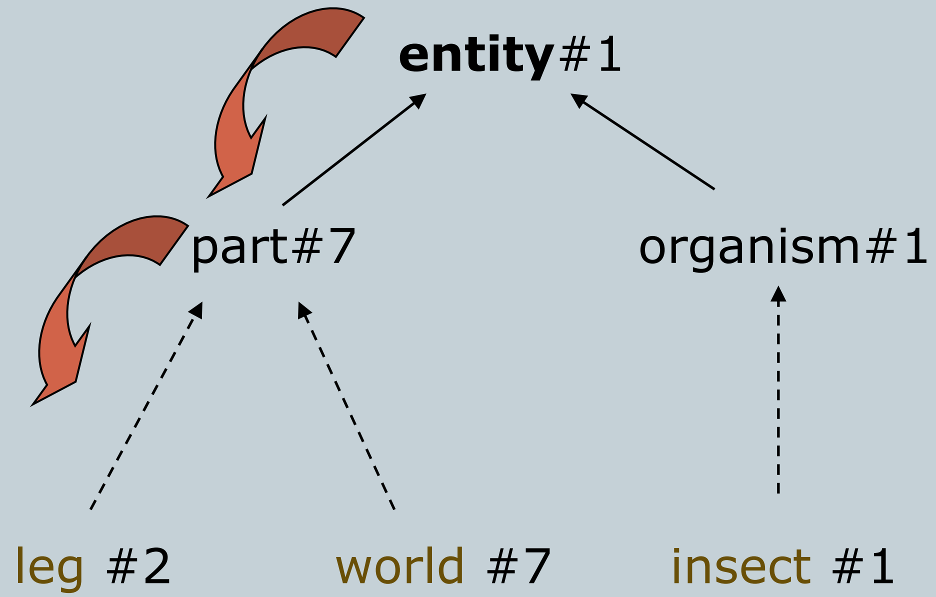

A Stochastic Model for Semantic Parsing

Experimental results

Relation	P	R	F	Acc	Total	Base-F	Base-Acc
Cause-Effect	69.5	**100**	**82**	**77.5**	80	67.8	51.2
Instrument-Agency	**68.2**	78.9	73.2	**71.8**	78	65.5	51.3
Product-Producer	84.5	79	81.7	**76.3**	93	80	66.7
Origin-Entity	86.4	52.8	65.5	**75.3**	81	61.5	55.6
Theme-Tool	85.7	41.4	**55.8**	**73.2**	71	58	59.2
Part-Whole	70.8	65.4	68	**77.8**	72	53.1	63.9
Content-Container	**93.1**	71.1	80.6	**82.4**	74	67.9	51.4
Average	79.7	71.1	80.6	**82.4**	74	67.9	51.4

A Stochastic Model for Semantic Parsing

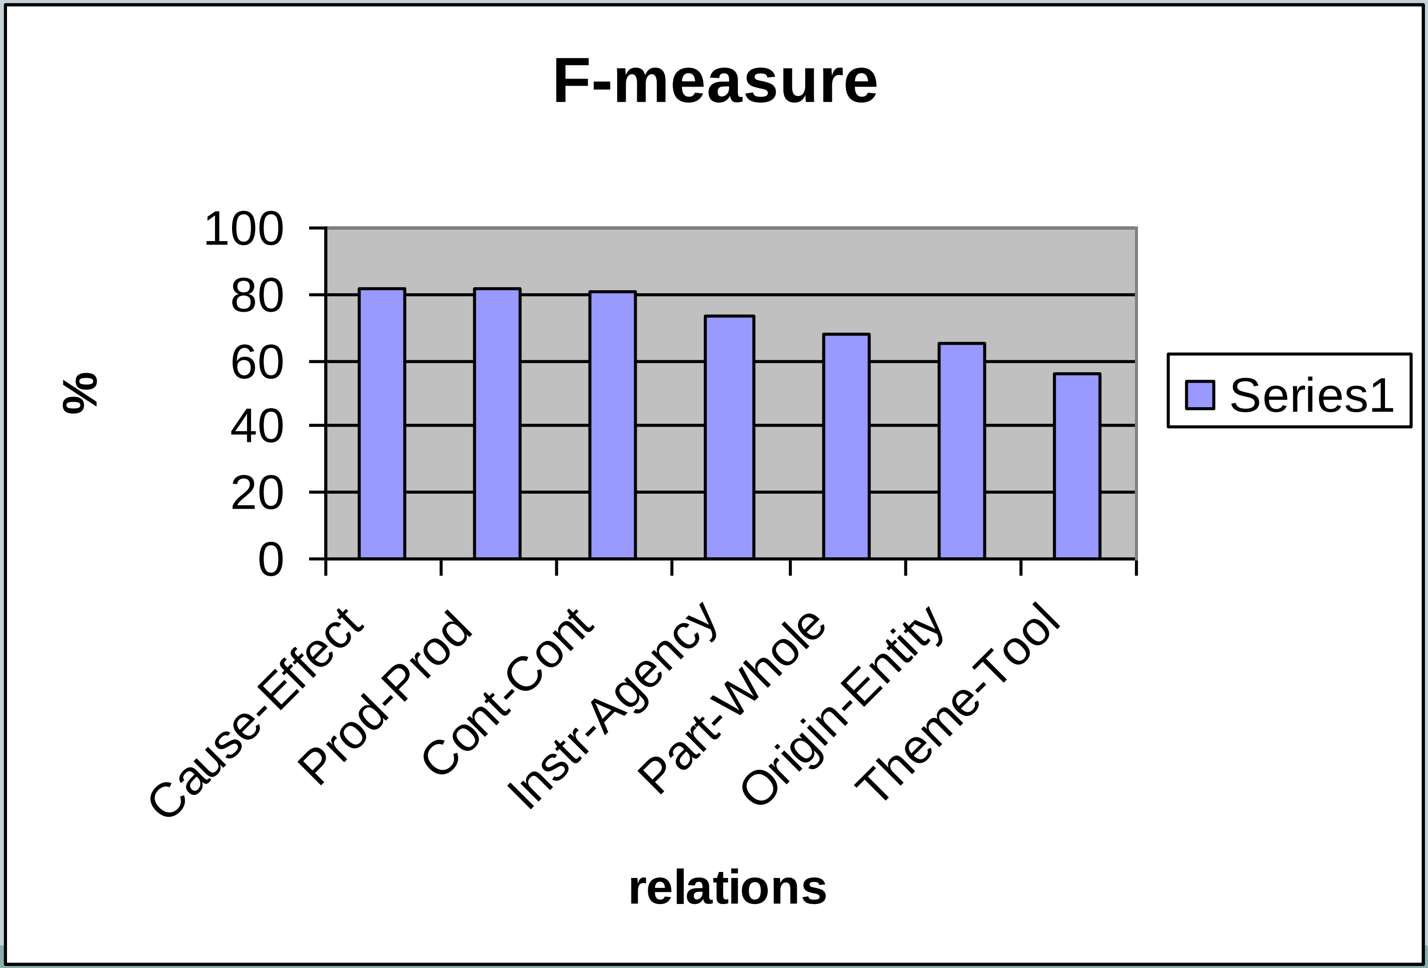

Application: Text-to-Scene Generation (1)

- Idea: Given a snipet of text, generate an image that is a faithful representation of that text

- Challenges:
  - Non-visual words (e.g., abstract words:policy, government, feeling)
  - Some knowledge needed for pictures is not explicitly stated in text, but inferred

- State of the art systems:
  - Fairy tales
    - WordsEye (Coyle and Sproat 2001): 3D objects (positions, color, texture, etc.); WordNet; (http://www.wordseye.com/)
  - Car simulations (car insurance purposes)
    - CarSim (Dupuy at al. 2001)
    - (Girju et al. 2011)

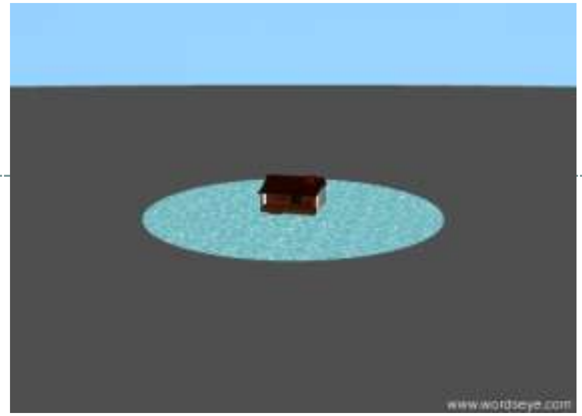

boat on the lake
vs. cabin on the lake

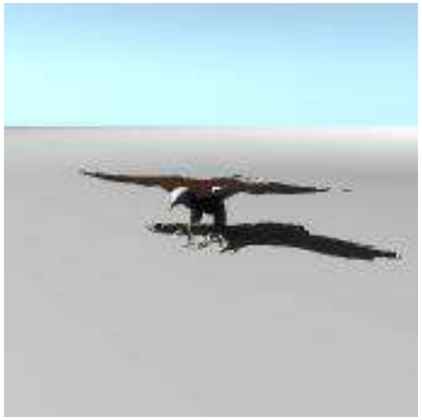

eagle in the nest
vs. eagle in the sky

flowers in a vase

Text-to-Scene Generation (3)

Car accident visualizations

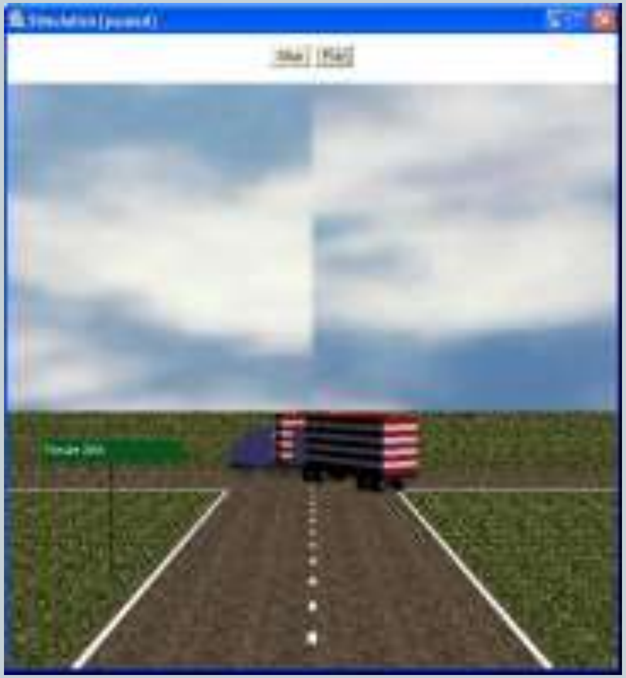

P2. Perspectives, Aspects and Sentiment

- in scientific literature
- in Israeli-Palestinian editorials
- cultural differences from travelers' experiences

Perspectives, Aspects and Sentiment

- TAM (Topic-Aspect Model):
  - Documents can be clustered along a number of dimensions: topics, sentiment/perspective/viewpoint
  - Discovers *topics* and *aspects*
  - Generates token assignment in both dimensions

Probabilistic Topic Models (1/2)

- Each word token associated with hidden "topic" variable
- Probabilistic approach to dimensionality reduction
- Useful for uncovering latent structures in text

- Basic formulation:
  - $P(w|d) = P(w|topic)\,P(topic|d)$

Probabilistic Topic Models (2/2)

- “Topics” are latent distributions over words (cluster of words)
- There are often other dimensions in which words could be clustered
  - Sentiment/perspective/viewpoint

- What if we want to model both?

Topic-Aspect Model (1/8)

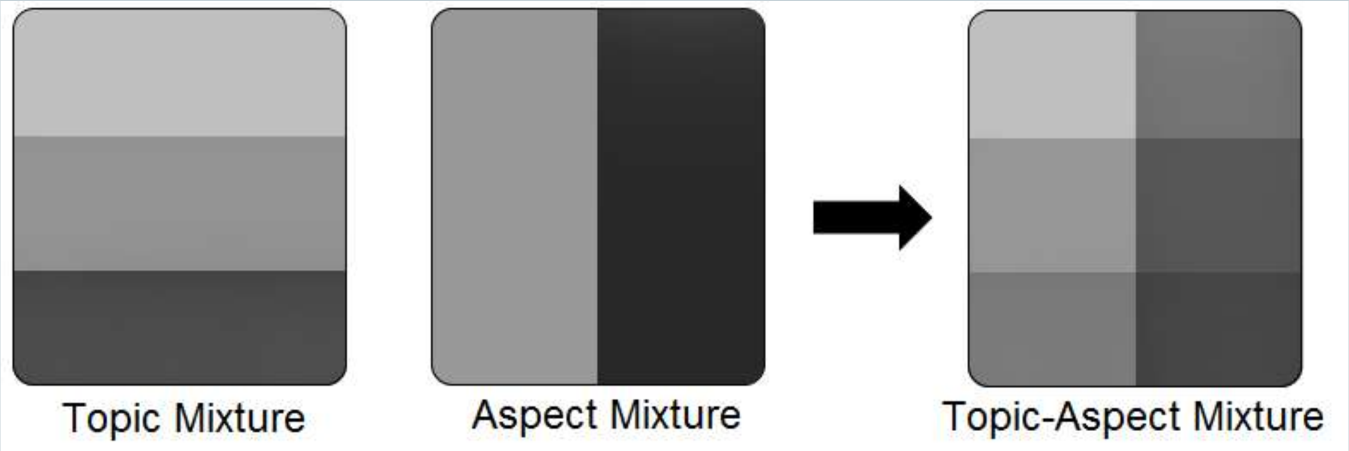

- Each document has
  - a multinomial topic mixture
  - a multinomial aspect mixture
- Words may depend on both!

Topic-Aspect Model (2/8)

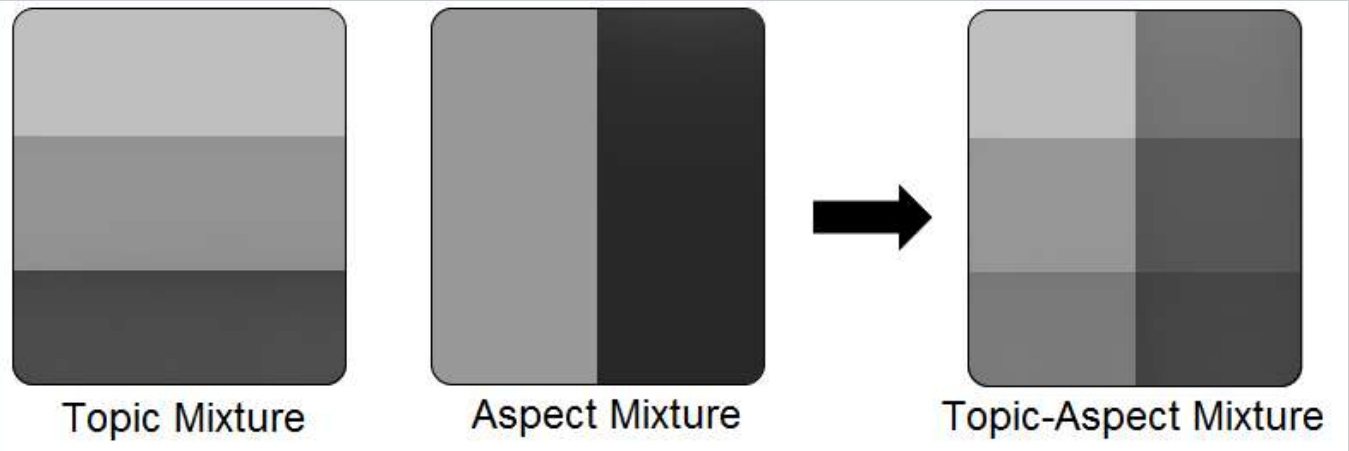

- Topic and aspect mixtures are drawn independently of one another
  - This differs from hierarchical topic models where one depends on the other
  - Can be thought of as two separate clustering dimensions

Topic-Aspect Model (3/8)

- Each word token also has 2 binary variables:
  - the "*l*evel" (background or topical)
    - denotes if the word depends on the topic or not
  - the "*route*" (neutral or aspectual)
    - denotes if the word depends on the aspect or not

- A word may depend on a **topic**, an **aspect**, **both**, or **neither**

Topic-Aspect Model (4/8)

Topic: SPEECH RECOGNITION
"Computational" Aspect

Route / Level	Background	Topical
Neutral	paper, new, present	speech, recognition
Aspectual	algorithm, model	markov, hmm, error

"Linguistic" Aspect

Route / Level	Background	Topical
Neutral	paper, new, present	speech, recognition
Aspectual	language, linguistic	prosody, intonation, tone

- A word may depend on a **topic**, an **aspect**, **both**, or **neither**

Topic-Aspect Model (5/8)

Topic: COMMUNICATION
"Computational" Aspect

Route / Level	Background	Topical
Neutral	paper, new, present	communication, interaction
Aspectual	algorithm, model	dialogue, system, user

"Linguistic" Aspect

Route / Level	Background	Topical
Neutral	paper, new, present	speech, recognition
Aspectual	language, linguistic	conversation, social

- A word may depend on a **topic**, an **aspect**, **both**, or **neither**

Topic-Aspect Model (6/8)

- Each token *i* is associated with: w_i, z_i, y_i, l_i, x_i
- Generative process for a document *d*:
  - Sample a topic z from $P(z|d)$
  - Sample an aspect y from $P(y|d)$
  - Sample a level l from $P(l|d)$
  - Sample a route x from $P(x|l,z)$
  - Sample a word w from either:
    - $P(w|l=0,x=0)$,
    - $P(w|z,l=1,x=0)$,
    - $P(w|y,l=0,x=1)$,
    - $P(w|z,y,l=1,x=1)$

Topic-Aspect Model (7/8)

- Distributions have Dirichlet/Beta priors
  - Latent Dirchlet Allocation framework
- Number of aspects and topics are user-supplied parameters
- Straightforward inference with Gibbs sampling

Topic-Aspect Model (8/8)

- Semi-supervised TAM when aspect label is known
- Two options:
  - Fix P($y|d$) = 1 for the correct aspect label and 0 otherwise
    - Behaves like ccLDA (Paul and Girju, 2009)
  - Define a prior for P($y|d$) to bias it toward the true label

Experiments (1/3)

- Datasets and settings:
  - 4,247 abstracts from the ACL Anthology (CL-Only)
    - Z = 25; Y = 2;
  - 594 articles from the Bitterlemons corpus (Lin et al., 2006)
    - a collection of editorials on the Israeli/Palestinian conflict (I-P dataset)
    - Both unsupervised and semi-supervised setting (Z = 12)

Experiments (2/3)

- Example: **Computational Linguistics (CL-Only)**

Background
Neutral
paper
based
approach
information
present
language
new
using
model
analysis
different
problem
set
describes
context
work

Background
Aspect A
results
method
corpus
using
data
task
performance
learning
text
evaluation
methods
automatic
features
experiments
accuracy
algorithm

Topical		
Aspect A	**Neutral**	**Aspect B**
TOPIC 1		
similarity	semantic	ontology
patterns	relations	conceptual
clustering	lexical	verbs
words	relation	verb
classification	relationships	concepts
distributional	nouns	hierarchy
occurrence	categories	objects
TOPIC 2		
segmentation	discourse	temporal
text	relations	expressions
segment	events	tense
segments	event	theory
local	structure	aspect
coherence	descriptions	referring
cohesion	time	spatial

Background
Aspect B
natural
language
processing
structure
representation
semantic
linguistic
text
knowledge
framework
general
generation
form
computational
implemented
theory

Experiments (3/3)

- Example: **Israeli/Palestinian Conflict (I-P)**

Unsupervised

palestinian israeli israel military civilians attacks	
Aspect A	**Aspect B**
war	violence
public	palestinians
government	occupation
media	resistance
society	intifada
terrorist	violent
soldiers	non
incitement	force

Prior for P(aspect|d) for true label

state israel solution palestine palestinian states borders	
Israeli	**Palestinian**
jewish	palestinians
arab	return
israeli	right
jews	refugees
population	problem
jordan	refugee
west	rights
south	resolution

Evaluation

- Cluster coherence
  - 5 human annotators
  - Compare against ccLDA and LDA (Z=25)
    - TAM clusters are as coherent as other established models
  - "word intrusion" method (Chang et al., 2009)

CL-LING

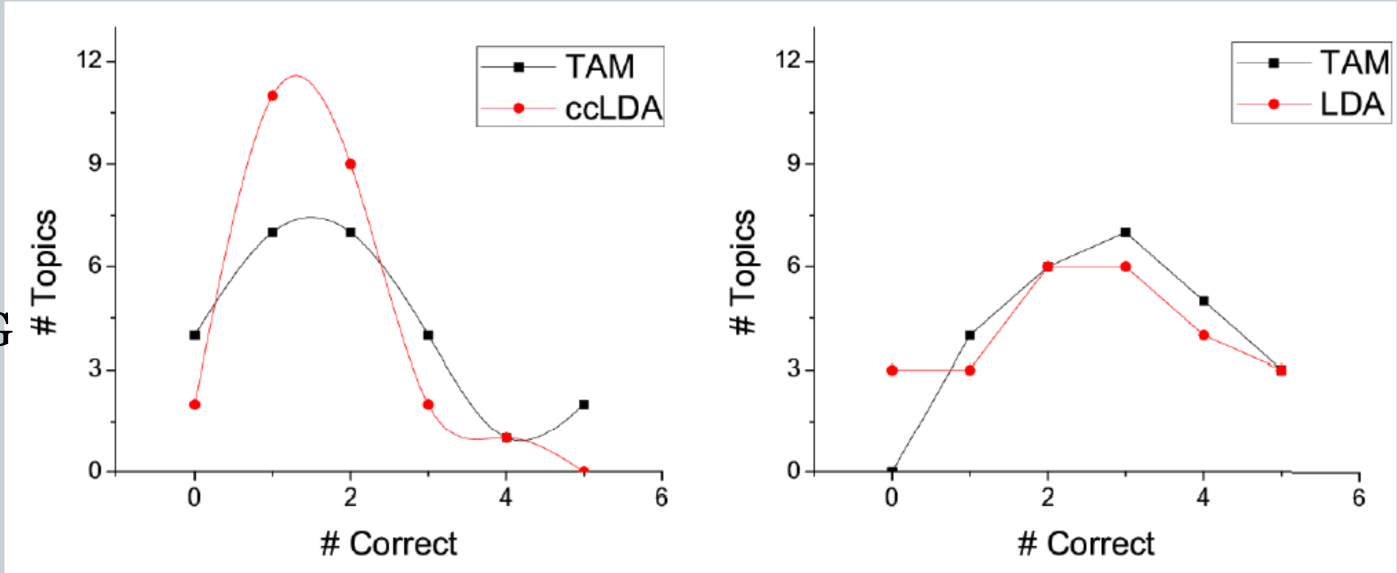

CL-ONLY

Other Applications: Social Media
(Paul and Girju 2009)

- Cultural differences in tourists' forums: Topic of 'weather' (tourist perspective)

(lonelyplanet.com)

weather time doing rain summer month high days thanks		
UK	**India**	**Singapore**
wind	leh	hot
waterproof	monsoon	humid
ending	road	heat
rolling	manali	degree
walkers	ladkh	equator
rochdale	trekking	sweat
layers	trek	bring
snow	season	rain
footwear	rains	umbrella
ankle	monsoons	

Other Applications: Social Media

- Cultural differences in tourists' forums: Topic of 'food' (blogcatalog.com)

Perspective of Locals			**Perspective of Tourists**		
food add chicken recipe cooking take rice recipes sugar soup			food eat restaurant tea cheap meal eating café drink		
UK	**India**	**Sing.**	**UK**	**India**	**Sing.**
food	recipe	coffee	chips	cooking	hawker
wine	recipes	cup	haggis	spices	satay
restaurant	powder	oil	fish	sick	stalls
coffee	indian	comments	respectability	flour	noodles
cheese	salt	fried	decent	tomato	roti
soup	tsp	add	veggie	batter	stall
eat	rice	rice	pudding	ate	seafood
english	masala	tea	photoblog	cook	malay
drink	oil	seafood	sausages	olive	sochester

P3. Causal Event Inference

- Identification of causal relations between events
- Applications: question answering, textual entailment

Causal Knowledge

- a pervasive feature of human language and theorising about the world
- important in text comprehension, entailment, automatic question answering and information retrieval (Goldman et al., 1999; Khoo et al., 2001; Girju, 2003)

.. a prerequisite to perform textual reasoning

Despite this, the specification of a satisfactory general analysis of causal relations has long proved difficult.

Our Approach

Qs:

How is causality encoded in language?

How is it perceived and used?

Models:

- unsupervised or semi-supervised
- knowledge poor

➡

- Baselines
- New insights

Experiment
(Riaz & Girju, ICSC 2010)

Hypothesis:

Scenario-specific events, contributing towards same objective in a domain, are likely to be dependent on each other and thus make good candidates for causal relationships.

Focus:

- Event sequences (Events: <[Subject] {verb} [Object]>)
- Scenario-specific events in news articles
  - Hurricane Katrina (447 articles, 189,840 word-tokens, 14,996 word-types)
  - Iraq war (556 articles, 304,481 word-tokens, 20,629 word-types)

Examples

Data set: Hurricane Katrina
Scenario: Hurricane Katrina disaster and damage.
Example: Katrina {hit} Florida late last week. Since Friday, Dallas-based Southwest airlines {canceled} more than 250 flights.
Causal Rel.: "Katrina {hit} Florida" → "Dallas-based Southwest airlines {canceled} more than 250 flights".
Type: Inter-sentential causal relationship.

Data set: Iraq War
Scenario: US accusations and the UN inspections.
Example: Bush {criticized} UN for {being ineffective}.
Causal Rel.: "UN {being ineffective}" → "Bush {criticized} UN".
Type: Intra-sentential causal relationship.

Notion of Causality (1)

- Selection criteria:
  - Need for a consistent set of annotation guidelines to capture our perception of causality (as expressed by language)
  - Be based on causal theories
    - provide the annotator with a relatively objective test (without relying on intuitions which will vary significantly from annotator to annotator)
  - The test should also be easy to perform mentally (w/o detailed philosophical knowledge about causality)
    - Subjectivity would certainly arise in cases where the annotator is unaware of how certain things in the world work;
    - But not a problem (many people share more/less the same baggage of commonsense knowledge)

Notion of causality (2)

- annotation test for causality:

  (Answering yes to both would mean the two events are causally related):

  (i) Does event A occur before event B?

  (ii) Keeping constant as many other states of affairs of the world in the given text context as possible, does modifying event A entail predictably modifying event B?

- Causal relations in a broad sense: contingency discourse relations:
  - purpose, reason, explanation, argument-claim (Sanders et al,1992; Mann & Thompson, 1988)

Approach

- Unsupervised model:

1. Identify Topic-Specific Scenarios and their Events
   - Discover topic-specific scenarios
   - Identify scenario-specific events

2. Generate Appropriate Event-Pair Candidates
   - Group events
   - Identify frequent event pairs

3. Learn Causal Relationships
   - Identify causal dependency
   - Assign cause and effect roles

1. Identifying Topic-Specific Scenarios and their Events

- cluster the input sentences according to their probability distributions into topic-specific scenarios.
- Use topic models to cluster semantically related text units

  Pachinko Allocation Model (PAM) (Lei & McCallum, 2006)
  - Assume documents are represented by Directed Acyclic Graph (DAG)
  - Cluster text units
  - Also find relationships between clusters

Topic-Specific Scenarios Discovered for the Iraq war

(C1) “War effects-economic progress in Iraq and side effects on the world’s economy”,
(C2) “US accusations and the UN inspection,
(C3) “Pre-war: War strategies and planning”.

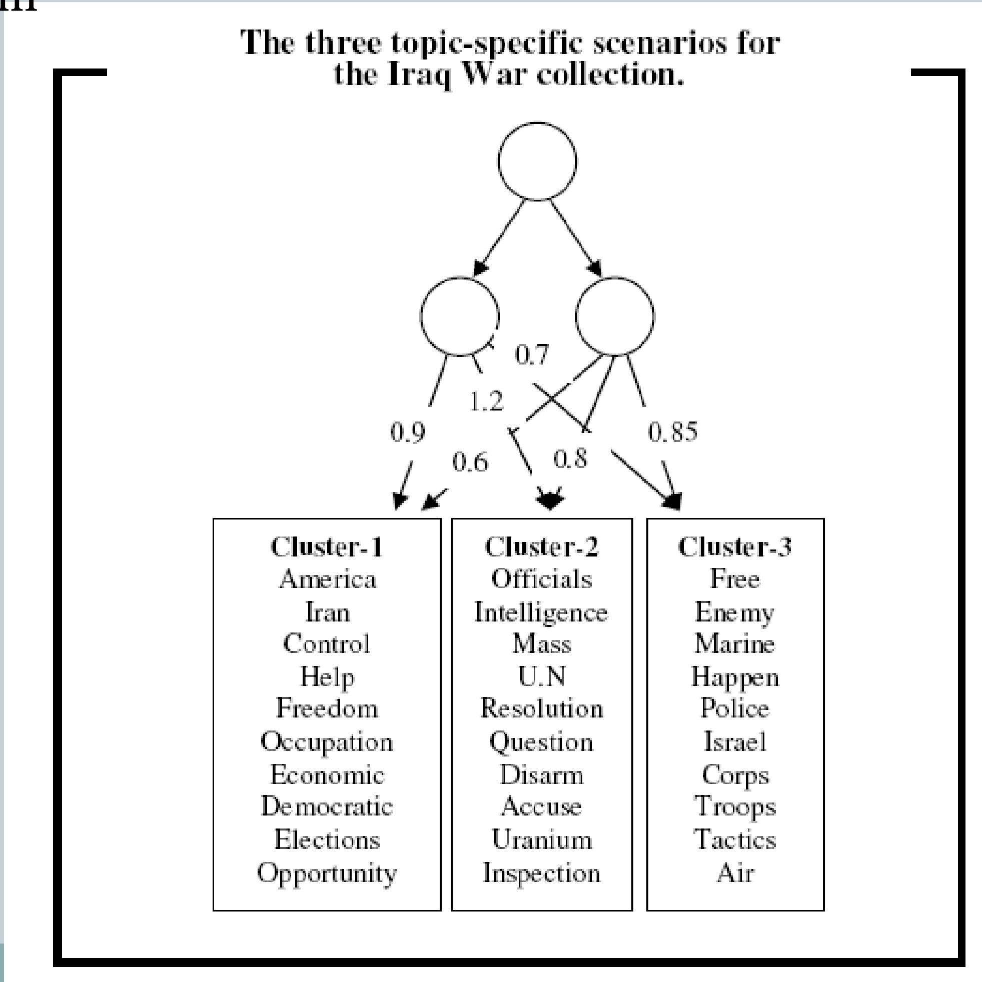

Recovering Sentences and their events from Discovered Scenarios

- Representation:
  - each scenario cluster *v* is a vector of words and each word's weight is its probability of assignment to scenario *v* cf. PAM.
  - each sentence *s* is a vector of words. Each word's weight is its probability of occurrence in a sentence.
  - Assign sentence *s* to scenario cluster *v* with which it has highest cosine similarity measure (N is vocabulary size)

$$Cosine - Sim(\vec{s}, \vec{v}) = \frac{\vec{s} \cdot \vec{v}}{\sqrt{\Sigma_{i=1}^{N} s_i^2}\sqrt{\Sigma_{i=1}^{N} v_i^2}}.$$

Extracting Events from Recovered Sentences

- use Semantic Role Labeler (Surdeanu et al, 2005)
  - A0 – agent and A1 – theme

Sentences along with their events (shown in italic) assigned to the scenarios identified for the Iraq war collection.

Scenario 1 – "War effects - economic progress in Iraq and side effects on the world's economy"

Event in context: <*Financial markets {wobble}*> as Iraq war unfolds.

Scenario 2 -- "US accusations and the UN inspections"

Event in context: <*Pentagon {fears} last-ditch Iraqi chemical attack*>.

Scenario 3 -- "Pre-war: War strategies and planning"

Event in context: If the <*Kurds join the Shiites*> in a general offensive against the Sunnis, <*the Sunnis will probably lose*>.

2. Generating Appropriate Event Pair Candidates

- Grouping Events
  - {"UN council suspects Iraq", "UN security council suspects Iraq"}
  - basic clustering approach
    - distance measure dependent on lexical similarity of two events with same verb lemma.
- Identifying Frequent Event Pairs
  - FP-Growth (Han et al, 2004) algorithm to collect frequent event pairs (a, b) that appear in at least *n* news articles.

FP-Growth

$D_i = \{e_1, e_2, .. e_n)$

Group Events and Documents

D1	$G_{1.suspect}$, $G_{2.kill}$, $G_{3.fall}$
D2	$G_{1.suspect}$, $G_{2.kill}$, $G_{4.go}$
D3	$G_{1.suspect}$, $G_{3.fall}$, $G_{5.inspect}$
D4	$G_{1.suspect}$, $G_{2.kill}$, $G_{3.fall}$

D_i is the *ith* document (news article). Replace events in D_i with their group ids.

FP-Growth:
- output: (Gi, Gj)
- min. support of 3

E.g.: (G1.suspect, G2.kill)
(G1.suspect, G2.fall)

(other pairs were rejected; E.g., (G1.suspect, G4.go)

3. Learning Causal Relationships

- (1) determine if the events of a frequent event pair (a,b) encode a contingency relation
- (2) identify the Cause and the Effect roles

- (1) Causal dependency:
  - Rank candidate (**a, b**) based on how strongly dependent the events are.
  - Condition: Cause can appear independently with other events, while Effect is expected to have a high likelihood of occurrence in the presence of the causing event (similar to Suppes, 1970)
  - The causal events can appear anywhere in or across documents → direct or indirect relationships.

  - Event-Control Dependency: ECD (a,b):

$$max\left(\frac{P(a,b)}{P(b)-P(a,b)+\gamma}\times\frac{P(a,b)}{max_i P(a,c_i)-P(a,b)+\gamma},\right.$$
$$\left.\frac{P(a,b)}{P(a)-P(a,b)+\gamma}\times\frac{P(a,b)}{max_i P(c_i,b)-P(a,b)+\gamma}\right)$$

- 2) Cause and Effect roles

Example

Pair-1 = ("US accused Iraq of developing chemical weapons" [92],
"UN inspected Iraqi scientists" [51]) [50]

Pair-2 = ("UN Security Council held an emergency session" [55],
"Security Council closed emergency session" [6]) [2]

Experiments and Evaluation

- Evaluating Scenarios
  - Three scenarios for each domain
  - Relatedness:
    - Annotation of top-50 words in each scenario (i.e. “YES” if a word is semantically similar to other words in top-50 list, otherwise “NO”)

Test	Data	C1	C2	C3
Relatedness	Katrina	66%	57%	65%
	Iraq	90%	83%	39.5%
Annotator-Agreement	Katrina	86%	94%	92%
	Iraq	80%	96%	86%

Evaluation of word relatedness and inter-annotator agreement for all three scenario clusters

Evaluating the Ranked Causal Relationships

Human Annotations

- Annotation guidelines : Manipulation Theory of Causality Test (Beamer & Girju, 2009)
- Observations:
  - the Iraq war instances are much more difficult to annotate.
  - cases where the annotators had to have some advanced domain knowledge in order to annotate the examples.

- Human perception task:
  - Does distance between events reduce strength of relationship between events?
  - HK(C1+C3) and IW(C1+C2): 80 event pairs each at distances 1 to 4 (20 examples)
  - The smaller the distance between two events, the more likely it is to be perceived as causal

Experiments and Evaluation (1)

- Causal Dependency Evaluation through interpolated precision-recall curve
  - Distance 1 and 2 only
  - top 100 ranked causal pair (a,b) examples from top two scenarios of both domains

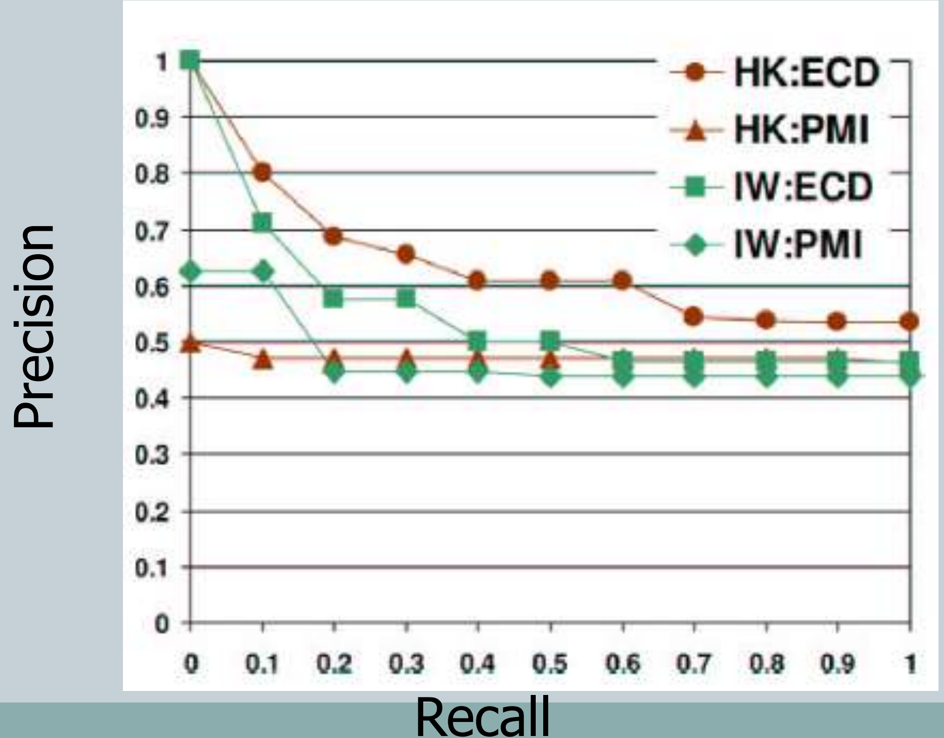

Experiments and Evaluation (2)

- Assignment of Cause and Effect Roles
  - Distance 1 and 2 only

Task	HK:ECD	HK:PMI	IW:ECD	IW:PMI
RA	63.4%	71.0%	72.5%	66.6%
CA-agreement	98%	93%	90%	85%
RA-agreement	100%	97%	95%	94%

Table 5: Roles Accuracy (RA) for all test sets. Roles accuracy = # of correct roles predicted/# of causal examples. CA-agreement and RA-agreement show the inter-annotator agreement on the causality annotation and causality roles annotation tasks.

Conclusions

- Unsupervised, knowledge-poor model
  - Baseline

- Improvements:
  - Better similarity measures
  - Better clustering models
  - Better analysis of the extracted causal relations:
    - Role of context vs. statistical tendencies
    - types of knowledge

Application: Social Reasoning

(Paul, Girju, Li, CONLL 2009; Girju ICWSM 2010; Girju & Paul, NLE 2011)

- Model people interactions:
  - **Direct observations** (explicitly stated)
  - **Inferences** of what they feel/think
  - **Predictions** of impending actions
  - **Decisions/suggestions** about subsequent interactions

Application: Social Reasoning

- When perceiving, explaining, judging human behavior, people differentiate between intentional / unintentional actions
  - This can set the course of social interactions

- Example:

Behavior perceived as <u>intentional</u>:	Behavior perceived as <u>unintentional</u>:
Critical remark → insult	Critical remark → constructive feedback
Collision in the driveway → provocation	Collision in the driveway → new friendship
Charming smile → hint of seduction	Charming smile → good mood
Negatively evaluated behavior	*Positively* evaluated behavior

Thank You