Topics for Today

- Evaluating research
  - A typical review form in NLP
  - Exercise (and a way to introduce some Cornell NLP research)
- Generating new research ideas

NAACL 2015 review form

- http://naacl.org/naacl-hlt-2015/review-form.html

Topics for Today

- Evaluating research
  - A typical review form in NLP
  - **Exercise – what makes a piece of research publishable?**
- Generating new research ideas

Evaluating Research

(via some of the recent research from my group)

The Research Topics

- Fine-grained opinion extraction
- Event extraction

Bishan Yang (PhD 2015)
postdoc at CMU

Fine-grained opinion extraction

- A type of Information Extraction
  - Text to structured information

Hillary Clinton on Tuesday defended Obamacare. She described some Republicans' legislative tactics to defund the landmark program as "deeply distressing" and "bad politics."
...

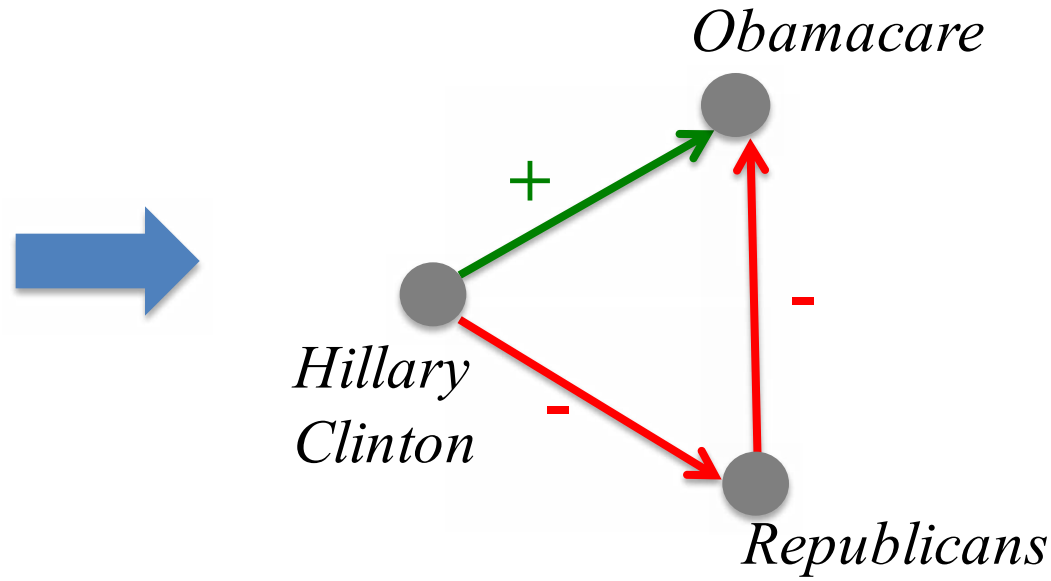

Paper should convince you that the task/problem tackled is important and/or interesting.

Opinion Frame

Hillary Clinton	*offered a defense*	*of*	*Obamacare* .
HOLDER	OPINION		TARGET

Opinion Frame

OPINION : *offered a defense*
HOLDER : *Hillary Clinton*
TARGET : *Obamacare*
POLARITY : positive
INTENSITY : medium

Event Frame

Hillary Clinton	*offered a defense*	*of*	*Obamacare* .
AGENT	EVENT		TARGET

Event Frame

EVENT	: *offered a defense*
AGENT	: *Hillary Clinton*
TARGET	: *Obamacare*
EVENT_TYPE	: defend

Extraction Subtasks

- **Entity extraction**
  - Identifying entities --- text spans that describe predefined objects or concepts, e.g., "*Hillary Clinton*", "*offered a defense*"

Extraction Subtasks

- **Entity extraction**
  - Identifying entities --- text spans that describe predefined objects or concepts, e.g., "*Hillary Clinton*", "*offered a defense*"
- **Relation extraction**
  - Identifying associations between entities, e.g., "*Hillary Clinton*" is an agent argument of "*offered a defense*"

Extraction Subtasks

- **Entity extraction**
  - Identifying entities --- text spans that describe predefined objects or concepts, e.g., "*Hillary Clinton*", "*offered a defense*"
- **Relation extraction**
  - Identifying associations between entities, e.g., "*Hillary Clinton*" is an agent argument of "*offered a defense*"
- **Attribute classification**
  - Assigning an attribute value to a given entity, e.g., "*offered a defense*" indicates positive sentiment

Extraction Subtasks

- **Entity extraction**
  - Identifying entities --- text spans that describe predefined objects or concepts, e.g., "*Hillary Clinton*", "*offered a defense*"
- **Relation extraction**
  - Identifying associations between entities, e.g., "*Hillary Clinton*" is an agent argument of "*offered a defense*"
- **Attribute classification**
  - Assigning an attribute value to a given entity, e.g., "*offered a defense*" indicates positive sentiment
- **Coreference resolution**
  - Identifying references of objects or concepts in text, e.g, "*She*" refers to "*Hillary Clinton*", and "*Hillary Clinton offered a defense of Obamacare*" and "*She defended Obamacare*" refer to the same event

Existing Solutions

- For opinions
  - **Opinion expression extraction** (e.g., Wiebe et al. (1999), Wiebe et al. (2005), Breck et al. (2007), Johansson and Moschitti (2010,2011))
  - **Holder extraction** (e.g., Bethard et al. (2004), Kim and Hovy (2004), Choi et al. (2005, 2006), Johansson and Moschitti (2010))
  - **Target extraction** (e.g., Wilson (2008), Stoyanov and Cardie (2008))
  - **Sentiment classification** (e.g., Wilson et al. (2009), Choi and Cardie (2008), Yessenalina and Cardie (2011), Socher et al. (2013))

Paper should make clear what is new and what is not new.

Existing Solutions

- For opinions
  - **Opinion expression extraction** (e.g., Wiebe et al. (1999), Wiebe et al. (2005), Breck et al. (2007), Johansson and Moschitti (2010,2011))
  - **Holder extraction** (e.g., Bethard et al. (2004), Kim and Hovy (2004), Choi et al. (2005, 2006), Johansson and Moschitti (2010))
  - **Target extraction** (e.g., Wilson (2008), Stoyanov and Cardie (2008))
  - **Sentiment classification** (e.g., Wilson et al. (2009), Choi and Cardie (2008), Yessenalina and Cardie (2011), Socher et al. (2013))
- For events
  - **Event trigger extraction** (e.g., Ahn (2006), Ji and Grishman (2008), Chen and Ji (2009))
  - **Event argument extraction** (e.g., Ahn (2006), Ji and Grishman (2008), Chen and Ji (2009))

Existing Solutions

- For opinions
  - **Opinion expression extraction** (e.g., Wiebe et al. (1999), Wiebe et al. (2005), Breck et al. (2007), Johansson and Moschitti (2010,2011))
  - **Hol**
    (20
  - and
    Cardie (et al. (2013))
- For events
  - **Event trigger extraction** (e.g., Ahn (2006), Ji and Grishman (2008), Chen and Ji (2009))
  - **Event argument extraction** (e.g., Ahn (2006), Ji and Grishman (2008), Chen and Ji (2009))

Limitations

① Addressing different extraction subtasks in isolation

② Lacking a discourse-level understanding of text

Paper should show that there are benefits to addressing these limitations.

Contribution:
Joint Inference for Opinion and Event Extraction

Key idea

Simultaneously considering different sources of low-level information and aggregating them across different parts of text

I: Joint inference across different extraction subtasks

- Joint opinion entity extraction and relation extraction (Yang and Cardie, ACL'13)
- Joint opinion expression extraction and attribute classification (Yang and Cardie, TACL'14)

II: Joint inference over multiple levels of contextual evidence

- Fine-grained sentiment analysis by leveraging intra- and inter-sentential cues (Yang and Cardie, ACL'14)
- Event coreference resolution within a document and across multiple documents (Yang, Cardie, and Frazier, TACL'15)

Outline

- Joint opinion entity and relation extraction
- Joint opinion expression extraction and attribute classification
- Discourse-aware fine-grained sentiment analysis
- Within- and cross-document event coreference resolution
- Conclusion & Future Work

Paper should make its hypothesis/hypotheses clear.
Nice if there are compelling examples to support them.

Opinion Entity and Relation Extraction

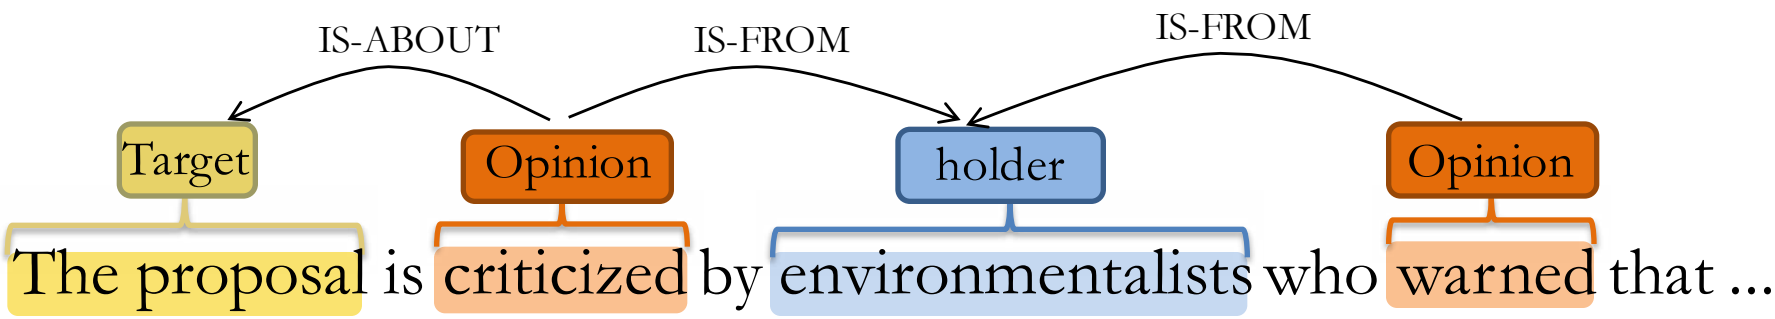

Definitions

Opinion expression: indicates an opinion, belief, emotion, evaluation,…
Opinion holder: specifies who holds the opinion
Opinion target: specifies the target or topic of the opinion

Related Work: Pipeline Approaches

- First identify opinion expressions and then identify the holders/targets of each opinion expression
  - Use grammatical rules or feature-based classification (Hu and Liu (2004), Kim and Hovy (2006), Kobayashi et al. (2007), Wu et al. (2009), Jakob and Gurevych (2010))

Error Propagation

The proposal is criticized by environmentalists who warned that ...

Opinion: 0.38
Source: 0.07
Target: 0.13
None: 0.42

Error Propagation

IS-ABOUT: 0.37

IS-FROM: 0.72

The proposal is criticized by environmentalists who warned that ...

Opinion: 0.01
Source: 0.71
Target: 0.25
None: 0.03

Opinion: 0.38
Source: 0.07
Target: 0.13
~~None: 0.42~~

Related Work: Joint Inference

- Named entity/relation classification (e.g., Roth and Yih (2004, 2007))
  - Assume entities are noun phrases
- Semantic role labeling (e.g., Punyakanok et al. (2008), Srikumar and Roth (2011), Das et al (2012))
  - Assume predicates are given
- Opinion holder extraction (Choi et al. (2006))
  - Assume only one type of opinion arguments; cannot handle missing arguments

In Bishan's case, there had already been work on this specific problem and variations of it AND a reasonable approach had been previously introduced.

Our Approach

- A new ILP formulation for joint opinion entity and relation extraction
  - Handle multiple types of opinion arguments
  - Handle missing arguments (implicit relations)

Joint Inference: an ILP formulation

$$\max_{x,u,v} \lambda \boxed{\sum_{i \in S} \sum_{z} f_{iz} x_{iz}} + (1-\lambda) \boxed{\sum_{k} \sum_{i \in O} \left(\sum_{j \in A_k} r_{ij} u_{ij} + r_{i\varnothing} v_{i\varnothing} \right)}$$

Opinion entity extractor

Opinion relation extractor

Comparison to pipeline approaches

	Opinion	Source	Target
CRF+Adj	73.31	55.56	58.97
CRF+Syn	73.31	44.29	53.28
CRF+RC	61.62	34.97	53.26
Joint-Model	**74.35**	**64.92**	**66.73**

Opinion entity extraction
(F1 using overlap matching)

	IS-ABOUT	IS-FROM
CRF+Adj	49.55	52.23
CRF+Syn	41.25	49.74
CRF+RC	32.28	50.00
Joint-Model	**57.04**	**61.63**

Opinion relation extraction
(F1 using overlap matching)

Make sure that the paper makes comparisons to the right baseline system(s).

Comparison to partially-joint approaches

	IS-ABOUT	**IS-FROM**
ILP-w/o-ENTITY	44.38	50.63
ILP-w-SINGLE-RC (Choi et al., 2006 + Implicit)	55.68	58.78
ILP-w/o-IMPLICIT-RC (Choi et al., 2006 + Target)	51.97	60.32
Joint-Model	**57.04**	**61.63**

Opinion relation extraction (F1)

Outline

- Joint opinion entity and relation extraction
- Joint opinion expression extraction and attribute classification
- Discourse-aware fine-grained sentiment analysis
- Within- and cross-document event coreference resolution
- Conclusion & Future Work

Event Coreference Resolution

- Goal: extract event mentions from text and group them within a document and across multiple documents

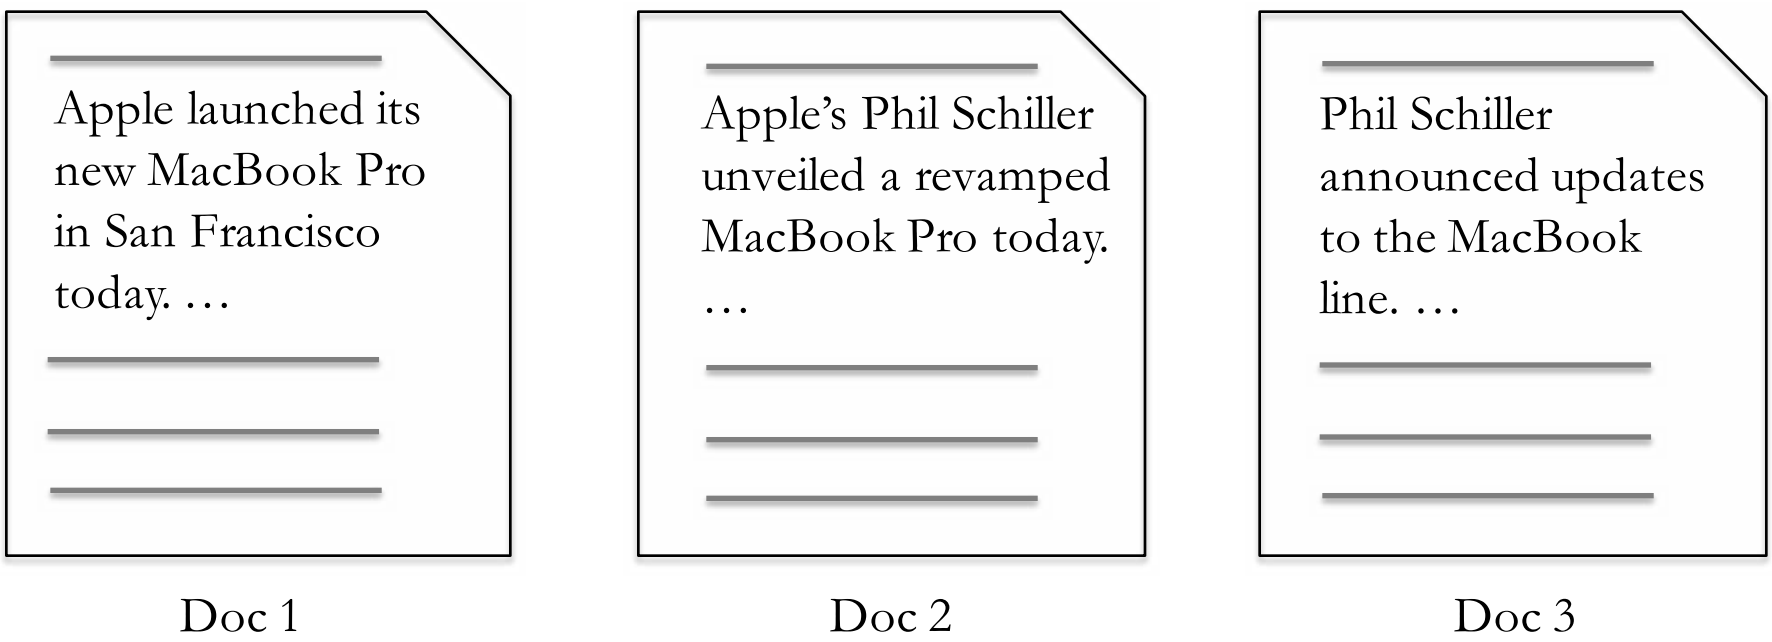

Event Frame Extraction

Apple	*launches*	*its*	*new MacBook Pro*	*in*	*San Francisco*	*today* .
PART	ACTION		PARTICIPANT		LOCATION	TIME

Definitions

ACTION	: what happens in the event
PARTICIPANT	: who or what is involved
LOCATION	: where the event happens
TIME	: when the event happens

Event Frame Extraction

Apple	*launches*	*its*	*new MacBook Pro*	*in*	*San Francisco*	*today* .
PART	ACTION		PARTICIPANT		LOCATION	TIME

Definitions

ACTION	: what happens in the event
PARTICIPANT	: who or what is involved
LOCATION	: where the event happens
TIME	: when the event happens

Approach: adapt the opinion frame extractor to extract event frames using event-related features

Event Coreference

- Two event (action) mentions are **coreferent** if they refer to the same **actual event**

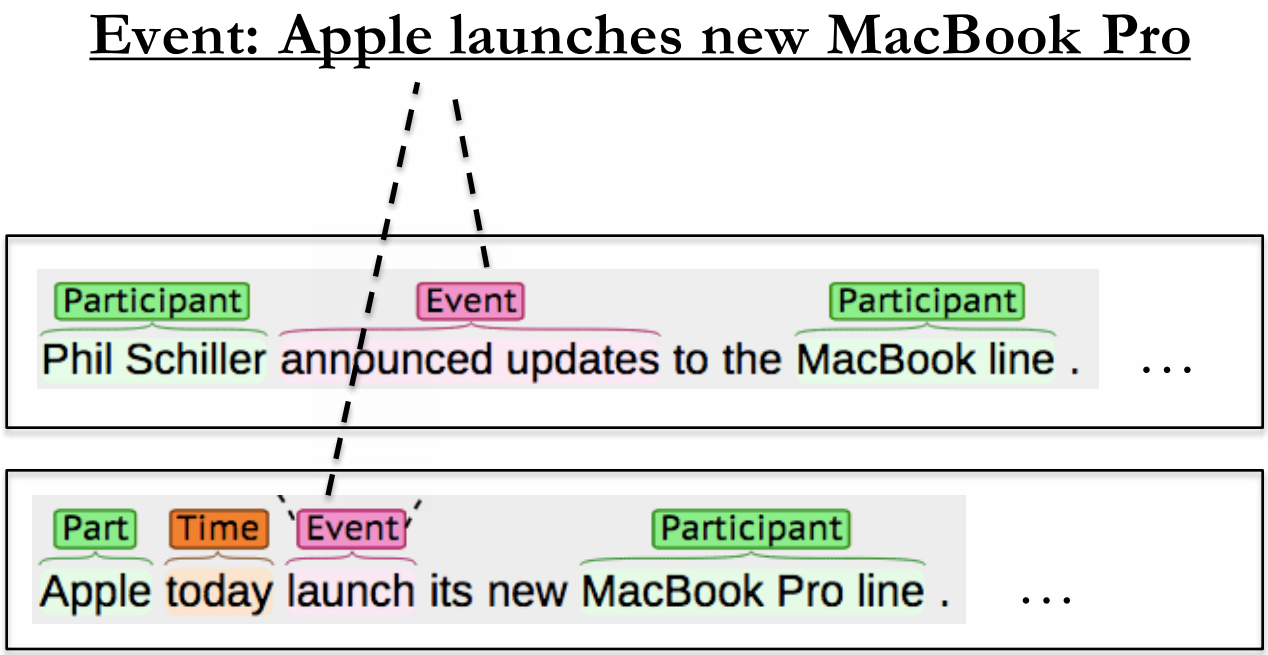

Existing Work

- Agglomerative clustering (deterministic) (Ahn, 2006; Chen et al., 2009)
  - Cannot capture global cluster structure
- Bayesian clustering (probabilistic) (Bejan and Harabagui, 2010; 2014)
  - Cannot capture similarities between mention pairs

MUCH less work has been done on this problem.

Our Approach

- Bayesian clustering with feature-rich similarity priors
  - An extension of the distance-dependent Chinese Restaurant Process (DDCRP) (Blei and Frazier, 2011)

Baselines to compare to should be pretty clear…

Coreference Results

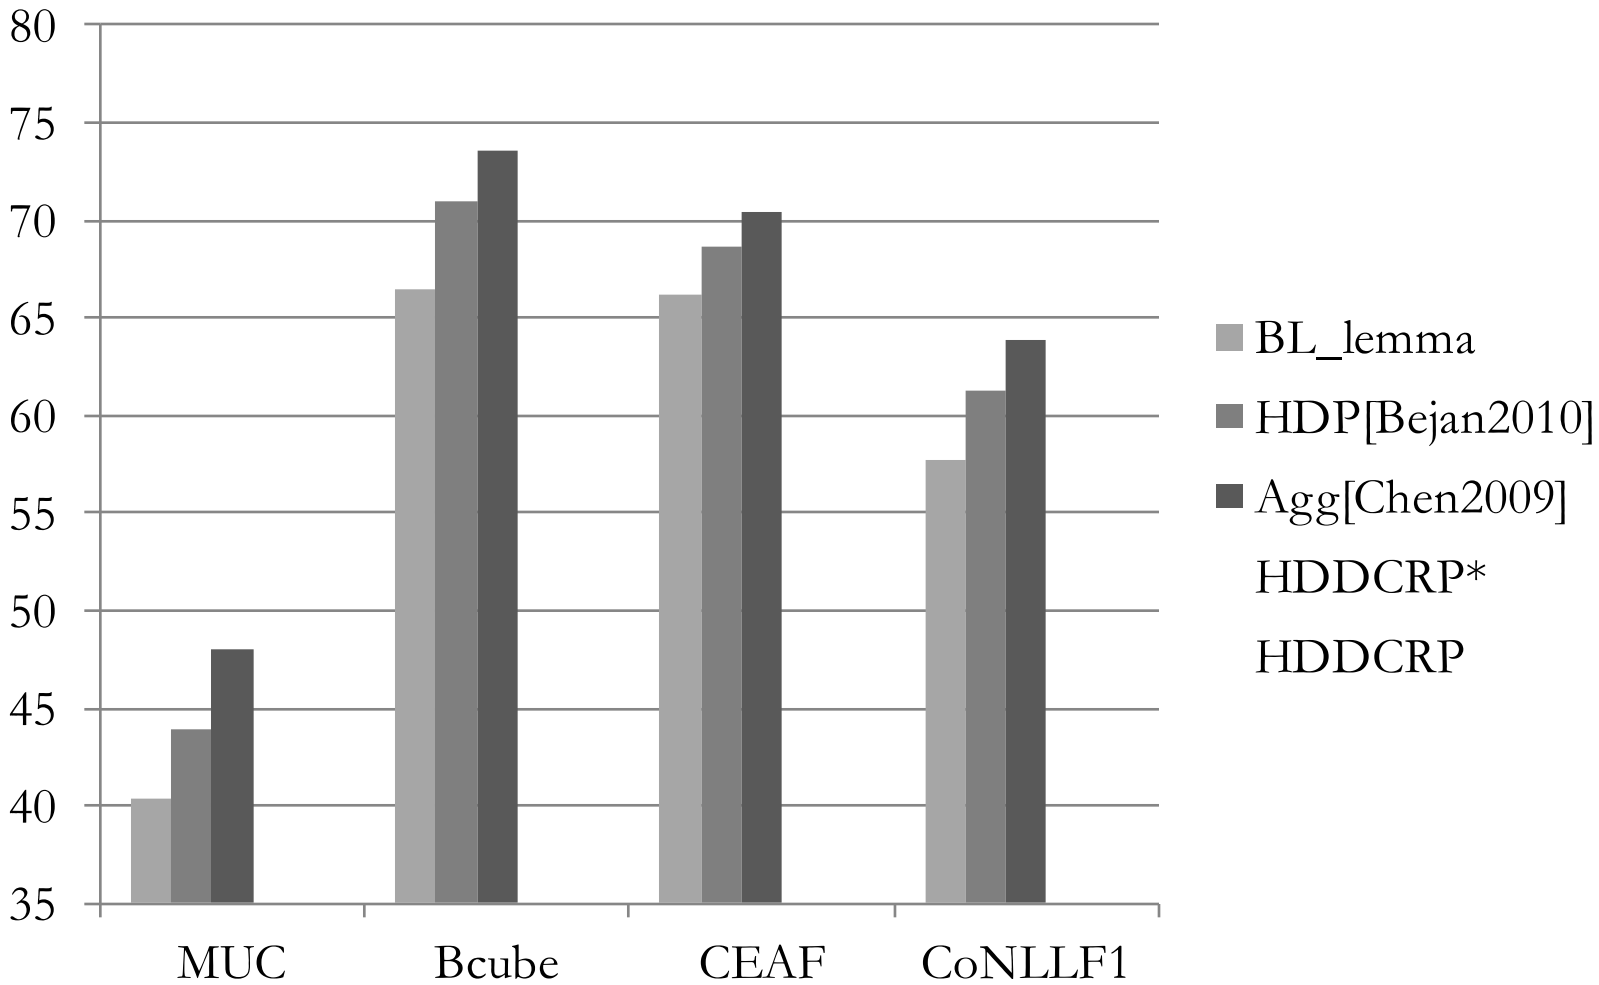

Within-document Coreference Evaluation (F1)

Coreference Results

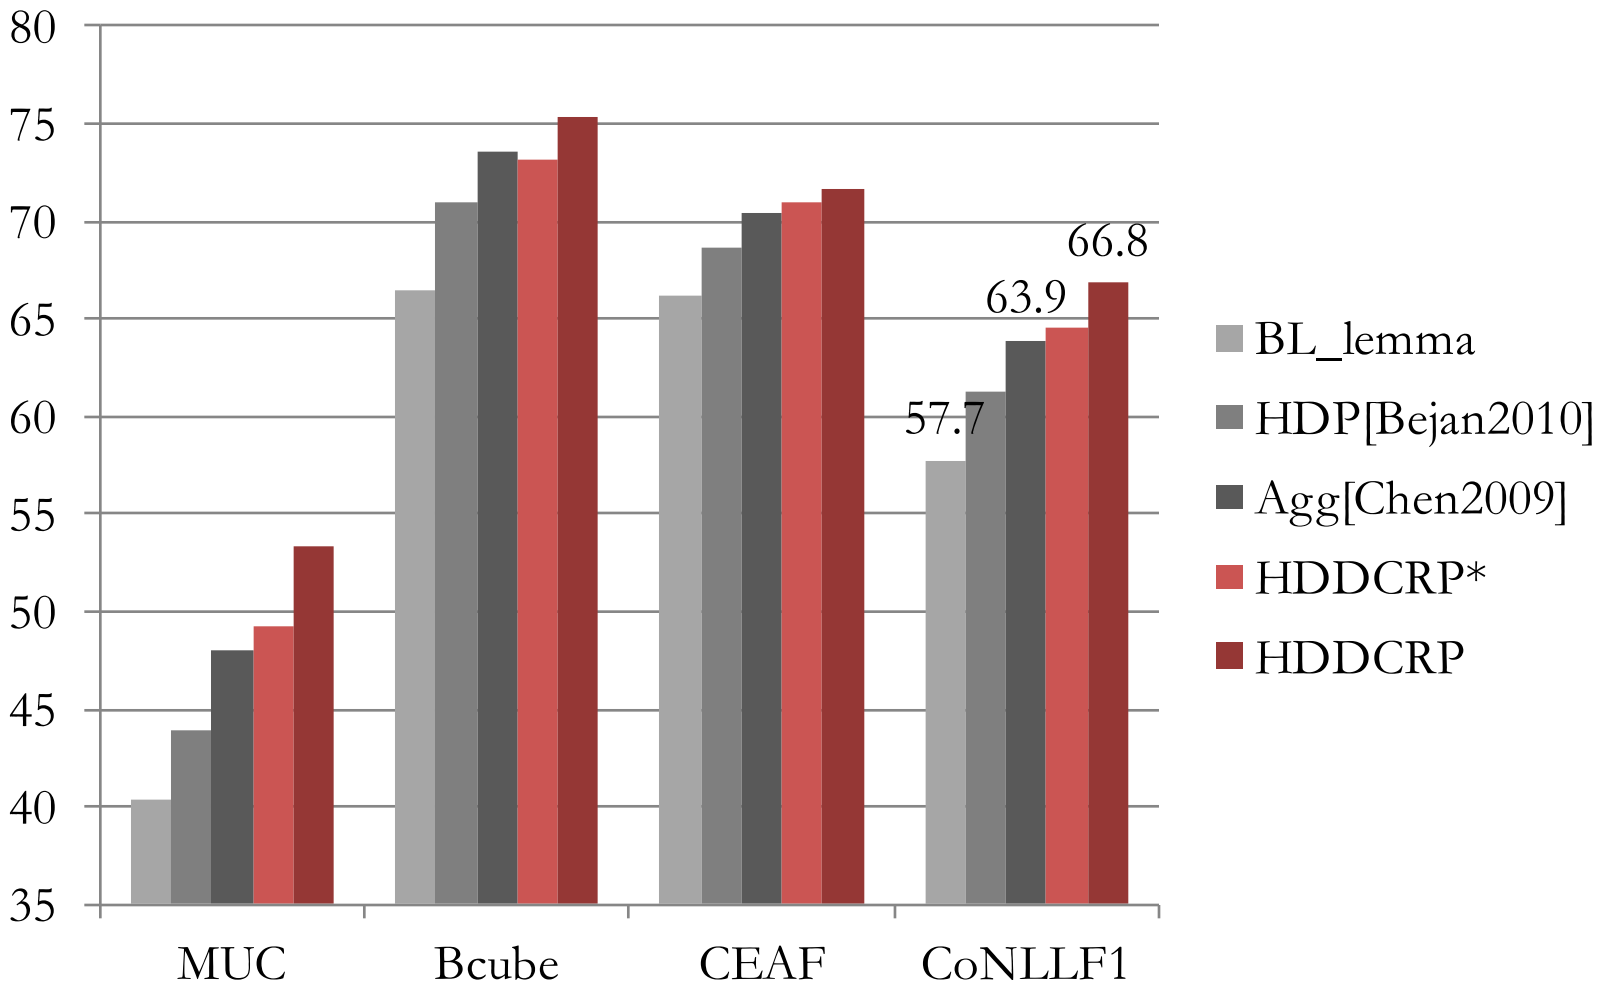

Within-document Coreference Evaluation (F1)

Coreference Results

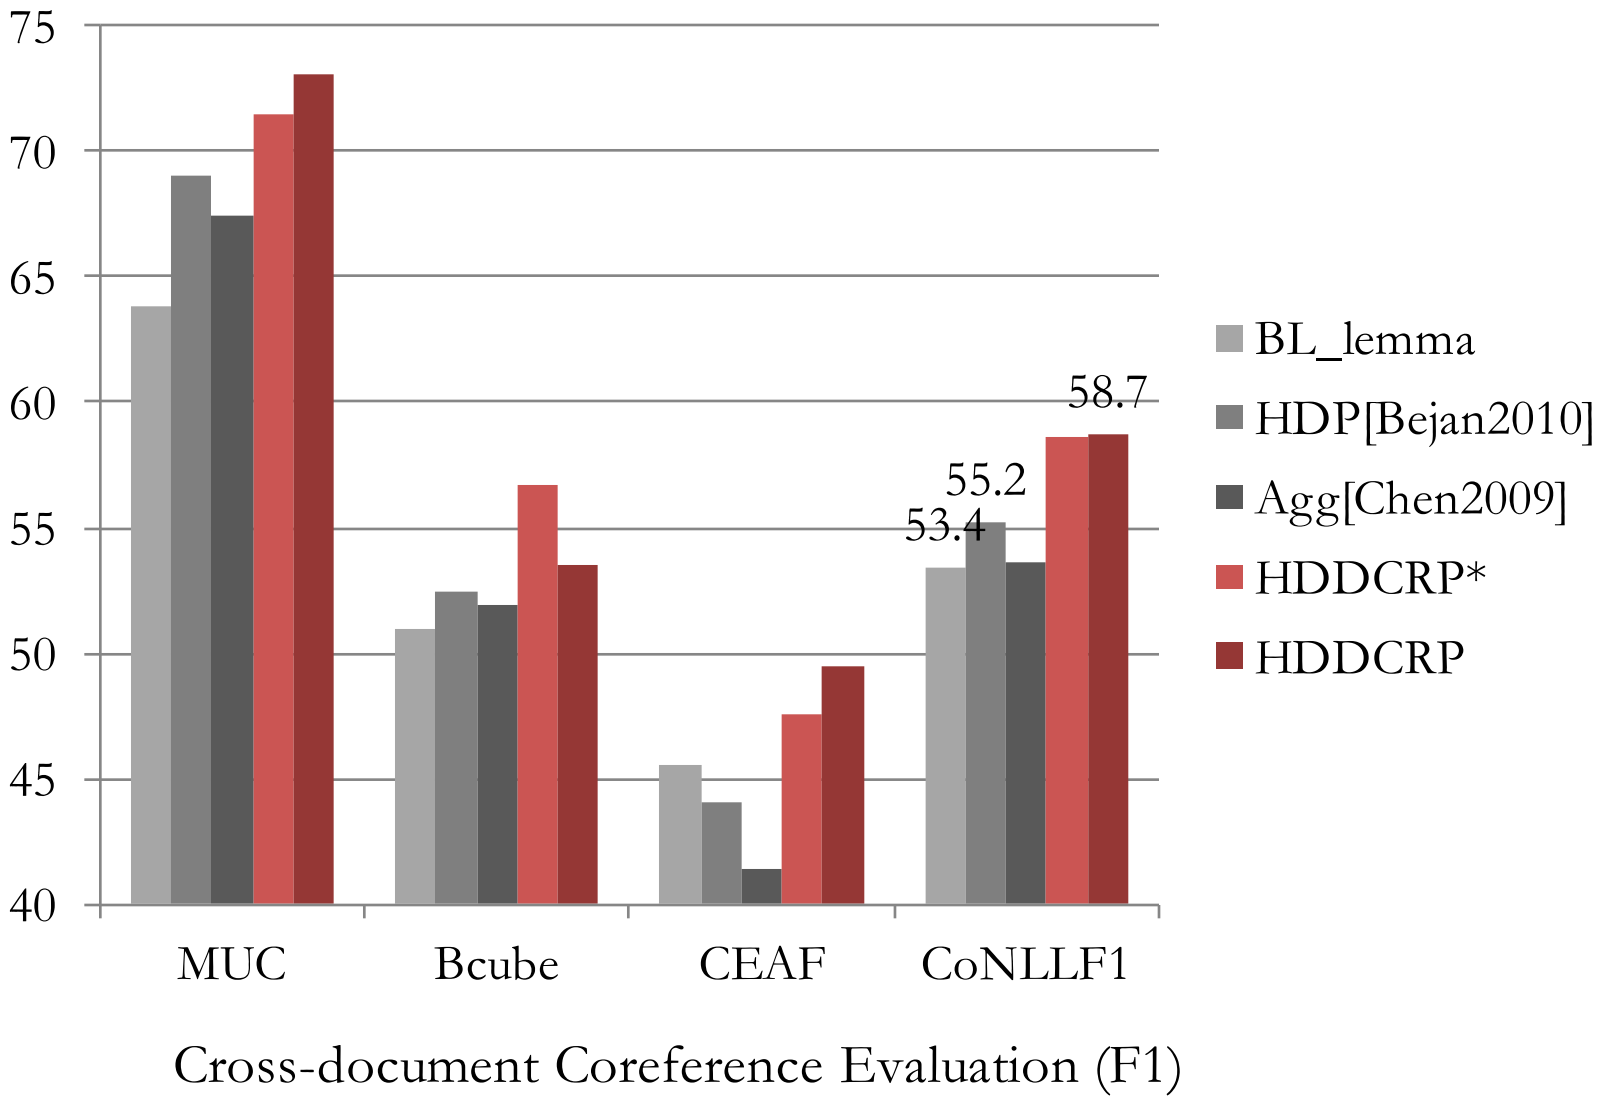

Topics for Today

- Evaluating research
  - A typical review form in NLP
  - Exercise
- **Generating new research ideas**

Questions to consider

- What is one of the strengths of the method proposed in the paper? How might it be used to solve a related problem? Examples are always good.
- What is one of the limitations of the paper's approach? Sketch out one or more possible solutions.
- Does the method described seem mature enough to use in real applications? Why or why not? What applications seem particularly amenable to this approach?

Questions to consider

- What good ideas does the problem formulation, the solution, the approach or the research method contain that could be applied elsewhere?
- What would be good follow-on projects and why? Sketch out how these might proceed.
- Are the paper's underlying assumptions valid? If not, explain why. How might you adapt the approach presented in the paper to fit your new assumptions.
- Are there other questions that might be investigated using the data sets(s)/corpora employed in the paper?