

Overview of the TAC2013 Knowledge Base Population **Sentiment Slot Filling Task**

Margaret Mitchell

Microsoft®
Research

JOHNS HOPKINS
UNIVERSITY



human language technology
center of excellence

Introduction

- New task this year
- **Sentiment** is defined as *a positive or negative emotion, evaluation, or judgement.*
- Explores the sentiment triple:
<sentiment holder, sentiment, sentiment target>
- We formalize this as:
{query entity, sentiment slot} → filler entity
- Entities: PER, ORG, GPE

Introduction

- Why is this task hard?
- *I love The Ravens!*
 - <writer, positive, Ravens>
- Naïvely:
 - Look for words like “love”, “hate”
 - Look for excited punctuation marks!!!
 - Look for emoticons :)

Introduction

- Sentiment is complex and nuanced
- *So **happy** that Kentucky lost to Kansas!!*
- *Had a **bad** time at the restaurant with **Mark**. :(*
*That place is the **worst**.*
- Linguistic variation in how people express ideas
- Often disagreement (at least in social media...)

Introduction (EMNLP 2013)

- Often disagreement (at least in social media...)

		Majority		
		Positive	Neutral	Negative
Minority	Positive	757	1249	130
	Neutral	707	2151	473
	Negative	129	726	452

Number of targeted sentiment instances where ≥ 2 of 3 annotators agreed on the polarity.

Challenges

- Discovering entities that are holders and targets of sentiment.
- Determining the polarity of the expressed sentiment.
- Determining which entities across documents are the same as the query entity (*this is its own difficult task*).
- **Bit of help:** Coref/NER using BBN's SERIF

Task Definition

- We are interested in:
 - Which entities hold sentiment towards another entity;
 - Which entities receive sentiment from another entity;
 - What the polarity of the expressed sentiment is
- Four list-valued slots:
 - **pos-towards**
 - **pos-from**
 - **neg-towards**
 - **neg-from**

Slots Definition

<sentiment holder, sentiment, sentiment target>

- **pos-towards:** query entity holds positive sentiment towards filler entity.
 - Fillers are *sentiment targets*.
- **pos-from:** query entity is a target of positive sentiment from filler entity.
 - Fillers are *sentiment holders*.
- **neg-towards:** query entity holds negative sentiment towards filler entity.
 - Fillers are *sentiment targets*.
- **neg-from:** query entity is a target of negative sentiment from filler entity.
 - Fillers are *sentiment holders*.

Further Guidelines

- Sentiment may be directed toward an entity based on direct evaluation of an entity.
 - e.g., *Kentucky doesn't like Mitch McConnell*
- Or may be directed to an entity based on actions that the entity took.
 - e.g., *Kentucky doesn't like Mitch McConnell's stance on gun control.*
- Given query with $\{Mitch\ McConnell, \mathbf{neg-from}\}$, filler would be holder of the sentiment, *Kentucky*.

Further Guidelines

- Post authors and bloggers may be used as query entities, or returned as filler entities.
 - If query, should be linked to KB or NIL.
- Complex sarcasm out of scope this year.

Filler Entities

- Entity strings must refer to distinct individuals.
 - If query includes *{Hillary Clinton, **pos-towards**}*, and system finds both “William Clinton” and “Bill Clinton”, just one, most informative should be returned.
- Entities should not be repeated as slot fillers for a single query.
 - Is possible that *Hillary Clinton* may feel **pos-towards** *William Jefferson Clinton* on many separate occasions; systems should only return one of these instances.

Provenance

- Return offsets for both query and filler entity.
- Sentences and clauses around the slot filler that provides justification for the extraction (at most two sentences).

Scoring and Assessment

- Pool responses, including manual LDC key.
- The slot filler in each non-NIL response is assessed as Correct, ineXact, or Wrong.
- Each correct response assigned to an equivalence class; credit for only one member of each class.

Scoring and Assessment

- Correct = total correct equivalence classes
- System = total non-NIL responses
- Reference = number equivalence classes for all slots

Then:

- **Precision** = Correct / System
- **Recall** = Correct / Reference
- **F1** = $2 * \text{Precision} * \text{Recall} / (\text{Precision} + \text{Recall})$

Participants and Systems

- Originally attracted 16 teams
- 3 teams submitted one or more runs
 - PRIS2013: Beijing University of Posts and Telecommunications
 - Columbia_NLP: Columbia University
 - CornPittMich: Cornell University, University of Pittsburgh, University of Michigan

Team ID	Run ID	Description
PRIS2013	1	CRF results with web metadata and forum metadata
	2	CRF results with forum metadata
	3	CRF results
Columbia_NLP	1	Opinion detection
	2	Opinion detection, subject / object filter
	3	High confidence opinion detection
	4	High confidence opinion detection, subject / object filter
	5	High confidence opinion detection, subject / object filter, subjectivity assumed
CornPittMich	1	Pipeline with opinion extraction followed by holder / target extraction

Participants and Systems

- Originally attracted 16 teams
- 3 teams submitted one or more runs
 - PRIS2013: Beijing University of Posts and Telecommunications
 - Columbia_NLP: Columbia University
 - CornPittMich: Cornell University, University of Pittsburgh, University of Michigan

Team ID	Run ID	Description
PRIS2013	1	CRF results with web metadata and forum metadata
	2	CRF results with forum metadata
	3	CRF results
Columbia_NLP	1	Opinion detection
	2	Opinion detection, subject / object filter
	3	High confidence opinion detection
	4	High confidence opinion detection, subject / object filter
	5	High confidence opinion detection, subject / object filter, subjectivity assumed
CornPittMich	1	Pipeline with opinion extraction followed by holder / target extraction

Participants and Systems

- Columbia_NLP and CornPittMich teams followed pipeline approach:
 - identify holders/targets
 - subjective expressions
 - sentiment polarity
 - developed this within-doc
- PRIS2013 followed relatively simpler pipeline:
 - identifying holders/targets
 - aggregate polarity over whole sentence

Participants and Systems

- Common approach was to use CRFs to identify sentiment holders and targets.
- PRIS2013 team used two models based on CRFs.
 - One to identify holders, one to identify targets.
- CornPittMich team incorporated the CRF/ILP-based system of Yang and Cardie (2013).
 - Identify subjective expressions, opinion targets, and opinion holders.

Participants and Systems

- All three teams used SERIF annotations for NER and coref.
- All teams additionally brought in Stanford CoreNLP tools for dependency parsing.
- All teams used some form of subjectivity or emotion lexicon.

Official Scores for SSF

Team ID	Run	Prec.	Rec.	F1
PRIS2013	1	9.15	20.24	12.60
	2	9.55	21.13	13.15
	3	5.50	12.17	7.58
Columbia_NLP	1	1.81	0.44	0.71
	2	1.83	0.44	0.71
	3	1.68	0.22	0.39
	4	1.72	0.22	0.39
	5	1.67	1.00	1.25
CornPittMich	1	10.00	0.77	1.44
LDC		70.01	75.66	72.73

- Official scores for Sentiment Slot Filling: Precision (Prec.), Recall (Rec.) and F-Score (F1) in %.

Scores for SSF with Some Leniency

Team ID	IGNOREOFFSETS			ANYDOC		
	P	R	F1	P	R	F1
PRIS2013	10.1	22.4	13.9	11.5	25.5	15.9
Columbia_NLP	1.9	1.1	1.4	2.6	1.6	2.0
CornPittMich	8.6	0.7	1.2	10.0	0.8	1.44

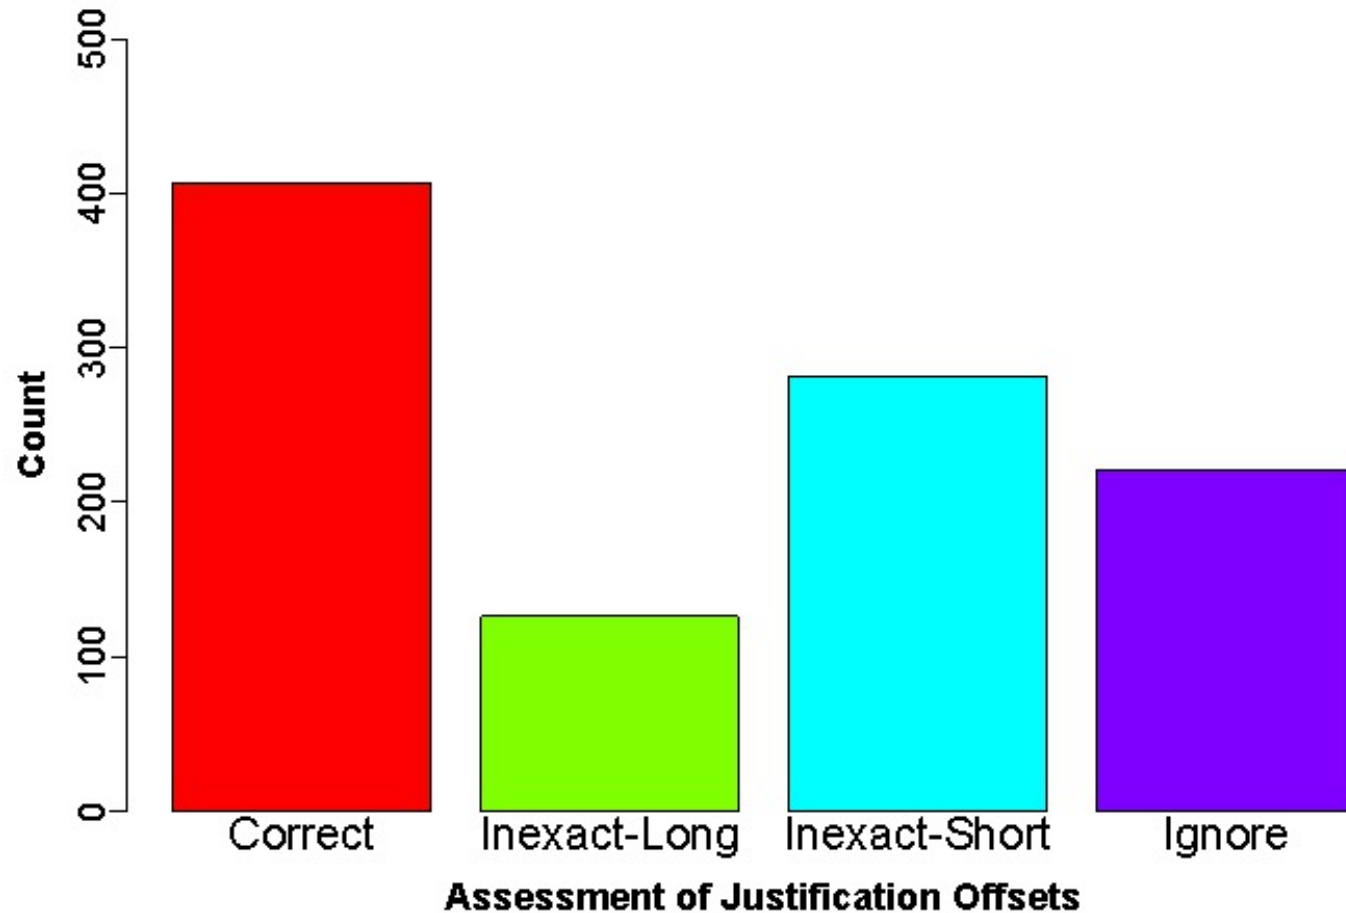
- Best team runs:
Precision (P), Recall (R) and F-Score (F1) in %.
- IGNORE-OFFSETS: justifications are considered correct if the correct document is reported.
- ANYDOC: justifications ignored, fillers marked correct based on string matching with gold fillers.

Correct Fillers Across Corpora

- Across teams, very few correct responses were drawn from the Web data.
- Discussion fora provided the richest source of correct slot fillers for this task.

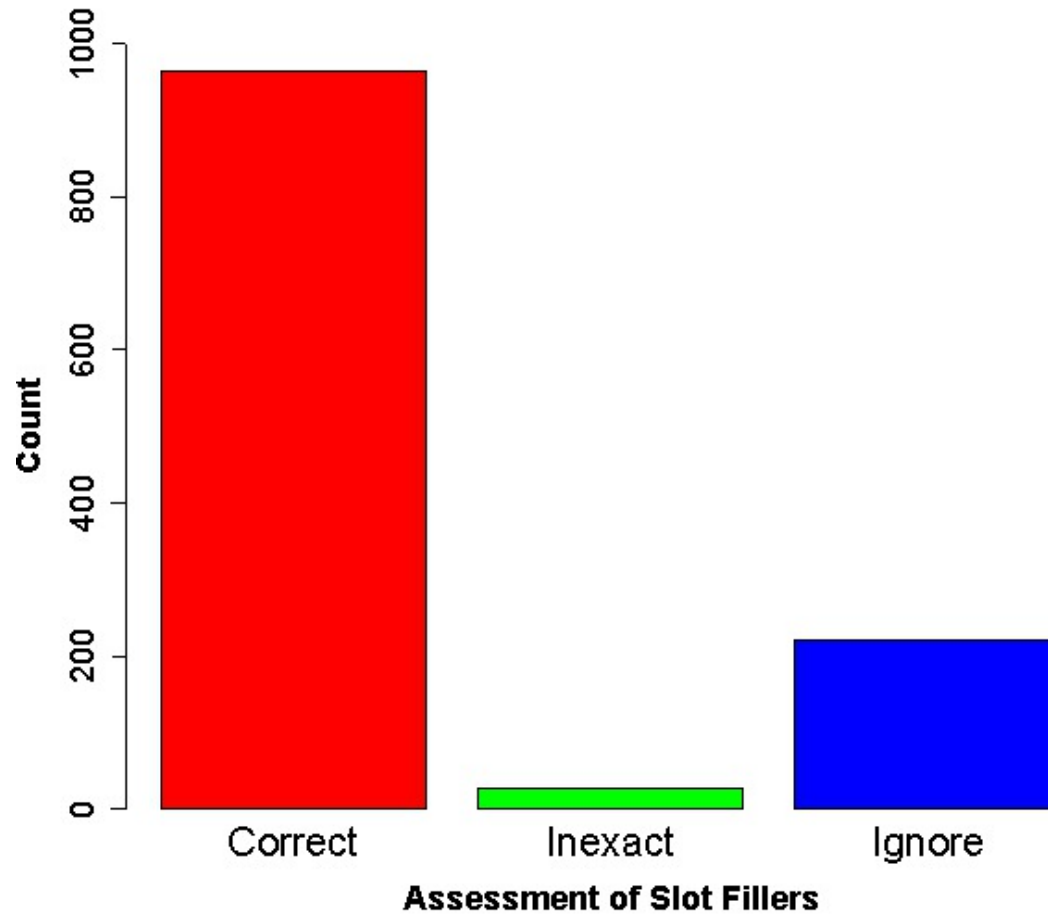
Slot	Data source			Total
	NEWS	WEB	FORA	
pos-towards	14	0	3	17
neg-towards	15	1	10	26
pos-from	11	0	72	83
neg-from	8	1	72	81
Total	48	2	157	207

Justification Assessment



- Excluding Wrong (4,124)

Slot Filler Assessment



- Excluding Wrong (3,947)

Example System Results

- PRIS:

slot: pos-from (positive sentiment from filler towards query)

query entity: Suzanne Collins, PER

filler: SMorriso (author in discussion fora)

justification: Also, Suzanne Collins' writing style was very stream of consciousness, imo.

- Not clear this is positive sentiment

filler: Rosemary B. Stimola (in newswire)

justification: Quite honestly, I knew from the very first paragraph I had a very gifted writer," says Stimola, who still represents Collins. "

slot: pos-towards (positive sentiment from query towards filler)

query entity: Avigdor Lieberman, PER

filler: Israel (in newswire)

justification: Israel sees "good chance" for dialogue with Palestine

- Not clear this is positive sentiment from Liberman towards Israel.

filler: Israeli (in newswire)

justification: Lieberman said Israel appreciates the traditionally good relations with Romania.

Example System Results

- Columbia_NLP:

slot: neg-towards (negative sentiment from query towards filler)

query: Cambodia, GPE

filler: Wen Jiabao

justification: Wen said, pledging to boost bilateral trade and implement infrastructure construction projects funded by China in Cambodia

- Does not show sentiment expressed by Cambodia
- May be sentiment expressed by Wen, though; but that would be positive

query: Erick Erickson, PER

filler: Mitch McConnell

justification: Erickson, the editor of the influential conservative blog RedState, is as hard on many Republicans and conservatives as he is on Democrats. He has accused Michael Steele, the chairman of the Republican National Committee, of playing the race card; suggested that RedState readers send toy balls to Sen. Mitch McConnell of Kentucky, the Republican leader

Future Directions

- Cross-document co-reference, entity linking
 - Ask participants to find fillers within a cluster of possible documents
- Simple approaches
 - Baseline system
- *Give* holders and targets; just ask for sentiment
- Dual Assessment?

Thanks!

- Thanks to: Ben Van Durme, Boyan Onyshkevych, Theresa Wilson, Mihai Surdenau, Hoa Trang Dang, Joe Ellis, Kira Griffit, Stephanie Strassel, and the rest of the KBP organizers
- Sara Rosenthal and Claire Cardie

