Welcome to TAC 2017!

- Please wear badges at all time while on NIST campus
- If you would like an airport shuttle or taxi to pick you up at NIST on Tuesday, sign up ASAP at the registration desk with your name and pick-up time. Otherwise, your taxi/shuttle will not be allowed past the security gate.
- This is a fully booked workshop. Please do not put personal items on the seat next to you; instead, use the space under or in front of your seat.

Text Analysis Conference
TAC 2017

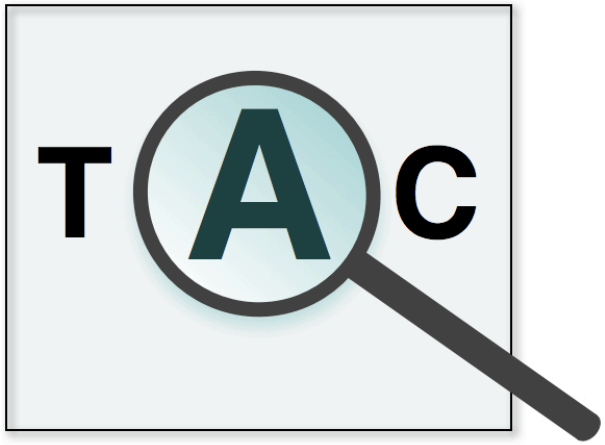

Sponsored by: NIST

Hoa Trang Dang

U.S. National Institute of Standards and Technology

Outline

- Intro to Text Analysis Conference (TAC)
- History of TAC tracks
- Overview of TAC 2017 tracks
- A word from our sponsor: Boyan Onyshkevych (DARPA)

TAC Goals

- To promote research in NLP based on large common test collections
- To improve evaluation methodologies and measures for NLP
- To build test collections that evolve to meet the evaluation needs of state-of-the-art NLP systems
- To increase communication among industry, academia, and government by creating an open forum for the exchange of research ideas
- To speed transfer of technology from research labs into commercial products

Features of TAC

- Component evaluations situated within context of end-user tasks (e.g., summarization, knowledge base population)
  - opportunity to test components in end-user tasks
- Test common techniques across tracks
- "Small" number of tracks
  - critical mass of participants per track
  - sufficient resources per track (data, annotation/assessing, technical support)
- Leverage shared resources across tracks (organizational infrastructure, data, annotation/assessing, tools)

Workshop

- "Working workshop" – audience participation encouraged
- Presenting work in progress
- Targeted audience is participants in the shared tasks and evaluations, objective is to improve system performance
  - Improve evaluation specifications and infrastructure
  - Discuss and investigate intriguing/unexpected evaluation results
  - Learn from other teams

TAC 2017 Track Participants – THANK YOU

- KBP Track coordinators
  - Cold Start KB/SF: Shahzad Rajput and NIST team
  - EDL: Heng Ji
  - Event: Marjorie Freedman and BBN/ISI team (Event Arguments); Teruko Mitamura and CMU team (Event Nuggets)
  - Belief and Sentiment: Owen Rambow and Columbia team
- ADR Track coordinators
  - Kirk Roberts, Dina Demner-Fushman, Joseph Tonning
- Linguistic resource providers:
  - Linguistic Data Consortium (Stephanie M. Strassel, Jeremy Getman, Jennifer Tracey, Zhiyi Song, ….)
- 55 Teams: 15 countries (25 USA, 14 China,….)

Ten Years of TAC Tracks

- Question Answering (2008)
- Recognizing Textual Entailment (2008-2011)
- Summarization (2008-2011, 2014)
- Knowledge Base Population (2009-2017)
  - DoD, (2009); DARPA Machine Reading (2010-2011), DEFT (2012-2017), AIDA (anticipated 2018)
- Adverse Drug Reaction Extraction from Drug Labels
  - FDA (2017)

ADR Extraction from Drug Labels (2017)

- Adverse reaction can be
  - Signs and symptoms
  - Changes in measures of critical body function (e.g., ECG)
  - Changes in laboratory parameters

- Task 1: Extract AdverseReactions and related entities (Severity, Factor, DrugClass, Negation, Animal).
- Task 2: Identify the relations between AdverseReactions and related entities (i.e., Negated, Hypothetical, Effect, and Equiv).
- Task 3: Identify the positive AdverseReaction entities in the labels.
- Task 4: Normalize positive AdverseReaction entity (strings) to MedDRA PTs.

Knowledge Base Population (2009 – 2017)

- Sponsored by US Department of Defense
- Goal: Populate a knowledge base (KB) with information about ***real world entities*** as found in a collection of source documents
- KB must be suitable for automatic downstream analytic tools; no human in the loop (contrast to KB as a visualization or browsing tool)
  - Input is unstructured text, output is structured KB
  - Follow a predefined schema for the KB (rather than OpenIE)
  - Confidence associated with each assertion whenever possible, to guide usage in downstream analytics
- Two use cases:
  - Augment an existing reference KB
  - Construct a KB from scratch (Cold Start KBP)

Knowledge Graph Representation of KB

Seymore Skinner

Springfield

per:cities_of_residence

Margaret Simpson

Entity

Location

Contact.
Meet

Entity

Marge Simpson

per:cities_of_residence

BOTTOMLESS PETE, NATURE'S CRUELEST MISTAKE

per:alternate_names

Homer Simpson

per:spouse

per:children

per:children

Springfield Elementary

Noncommited
Belief

per:schools_attended

Lisa Simpson

Negative
Sentiment

Bart Simpson

Difficult to evaluate KBP as a single task

- Wide range of capabilities required to construct a KB
- KB construction is a complex task, but open community tasks are usually small (suitable even for a single researcher)
- Barrier to entry even greater when require multi-lingual processing and cross-lingual fusion
- KB is a complex structure → single-point estimator for KB quality provides little diagnostics for failure analysis

TAC approach to KBP evaluation

- Decompose the KB construction task into smaller components
- Allow participation in single component tasks, and evaluate each component separately
- Incrementally increase difficulty of tasks, building infrastructure along the way; provide component-specific evaluation resources to allow component capabilities to mature and develop in their own way
- As technology matures, incorporate components into a real KB and evaluate as part of the KB

KBP tracks

- Component tasks
  - Entities: 2009-present
  - Relations ("Slot Filling"): 2009-present
  - Events: 2014-present
  - Sentiment: 2013-2014, 2016-present
  - Belief: 2016-present
- Composite KB construction task ("Cold Start")
  - Entities, Relations: 2012-2016
  - Entities, Relations, Events, Sentiment: 2017

KBP COMPONENTS AND TASKS

Entity Tasks: 2009 => 2016

- Input
  - A large set of raw documents in English, Chinese and Spanish
    - Genres include newswire, discussion forum
- Output
  - Document ID, offsets for mentions (including nested mentions)
  - Entity type: GPE, ORG, PER, LOC, FAC
  - Mention type: name, nominal
  - Reference KB link entity ID, or NIL cluster ID
  - Confidence value

- Entity Discovery and Linking (EDL) produces KB entity nodes from raw text, including all named and nominal mentions of each entity

Relations: SF slots derived from Wikipedia infobox

Person		Organization
per:alternate_names		org:alternate_names
per:date_of_birth	per:employee_or_member_of	org:political_religious_affiliation
per:age	per:religion	org:top_members_employees
per:country_of_birth	per:spouse	org:number_of_employees
per:stateorprovince_of_birth	per:children	org:members
per:city_of_birth	per:parents	org:member_of
per:date_of_death	per:siblings	org:subsidiaries
per:country_of_death	per:other_family	org:parents
per:stateorprovince_of_death	per:charges	org:founded_by
per:city_of_death		org:date_founded
per:cause_of_death		org:date_dissolved
per:countries_of_residence		org:country_of_headquarters
per:statesorprovinces_of_residence		org:stateorprovince_of_headquarters
per:cities_of_residence		org:city_of_headquarters
per:schools_attended		org:shareholders
per:title		org:website

Events

Event Label (Type.Subtype)	Role	Allowable ARG Entity/Filler Type
Conflict.Attack	Attacker	PER, ORG, GPE
	Instrument	WEA, VEH, COM
	Target	PER, GPE, ORG, VEH, FAC, WEA, COM
Conflict.Demonstrate	Entity	PER, ORG
Contact.Broadcast	Audience	PER, ORG, GPE
	Entity	PER, ORG, GPE
Contact.Contact	Entity	PER, ORG, GPE
Contact.Correspondence	Entity	PER, ORG, GPE
Contact.Meet	Entity	PER, ORG, GPE
Justice.Arrest-Jail	Agent	PER, ORG, GPE
	CRIME	CRIME
	Person	PER
Life.Die	Agent	PER, ORG, GPE
	Instrument	WEA, VEH, COM
	Victim	PER
Life.Injure	Agent	PER, ORG, GPE
	Instrument	WEA, VEH, COM
	Victim	PER
Manufacture.Artifact	Agent	PER, ORG, GPE
	Artifact	VEH, WEA, FAC, COM
	Instrument	WEA, VEH, COM
Movement.Transport-Artifact	Agent	PER, ORG, GPE
	Artifact	WEA, VEH, FAC, COM
	Destination	GPE, LOC, FAC
	Instrument	VEH, WEA
	Origin	GPE, LOC, FAC
Movement.Transport-Person	Agent	PER, ORG, GPE
Personnel.Elect	Agent	PER, ORG, GPE
	Person	PER
	Position	Title
Personnel.End-Position	Entity	ORG, GPE
	Person	PER
	Position	Title
Personnel.Start-Position	Entity	ORG, GPE
	Person	PER
	Position	Title
Transaction.Transaction	Beneficiary	PER, ORG, GPE
	Giver	PER, ORG, GPE
	Recipient	PER, ORG, GPE
Transaction.Transfer-Money	Beneficiary	PER, ORG, GPE
	Giver	PER, ORG, GPE
	Money	MONEY
	Recipient	PER, ORG, GPE
Transaction.Transfer-Ownership	Beneficiary	PER, ORG, GPE
	Giver	PER, ORG, GPE
	Recipient	PER, ORG, GPE
	Thing	VEH, WEA, FAC, ORG,COM

Event Nuggets, Arguments, and Linking

- Given:
  - Source documents
  - Event Taxonomy
- Event Nugget task:
  - Detection all mentions of events from the taxonomy, and corefer all mentions of the same event (within-doc)
- Event Argument task:
  - Extract instances of arguments that play a role in some event from the taxonomy, and link arguments for the same event (within-doc)
  - Link coreferential event frames across the corpus (2016)
  - Don't have to identify all mentions (nuggets) of the event

Belief and Sentiment (BeSt)

- Detect belief (Committed, Non-Committed, Reported) and sentiment (positive, negative), including source and target
- Sources are Entities (person, organization, geopolitical entity)
- Targets can be:
  - Entities: for sentiment ("Mary likes John")
  - Relations: for belief ("John believes Mary was born in Kenya") and sentiment ("John doesn't like that Mary was president")
  - Events: for belief ("John thought there might have been demonstrations supporting his election") and sentiment ("John loved the demonstrations from the animal rights group")
- Possible source entities and targets are given as input, BeSt system focuses on detecting belief/sentiment between them.

COMPOSITE KBP TASK

2012-2016 Cold Start KB Construction Task

- Goal: Build a KB from scratch, containing all attributes about all entities as found in a corpus
  - Entity Discovery and Linking system component identifies KB entities and all their NAM/NOM mentions
  - Slot Filling system component identifies entity attributes (fills in "slots" for the entity)
- Process one batch of 30-60K English documents
- Post-submission slot filling evaluation queries traverse KB starting from a single entity mention (entry point into the KB):
  - Hop-0: "Find all children of Marge Simpson"
  - Hop-1: "Find schools attended by each child of Marge Simpson"

KBP 2017

- Component KBP tasks and evaluations
  - EDL
  - Slot Filling
  - Event Nuggets, Event Argument Extraction and Linking
  - Belief and Sentiment
- Composite Cold Start KB Construction task
  - Systems construct KB from raw text. KB contains:
    - Entities
    - Relations (Slots)
    - Events
    - Sentiment towards entities
  - KB populated from English, Chinese, and Spanish (30K/30K/30K docs)
  - Confidence-aware metrics and multiple justifications

Trilingual Cold Start KB (2017)

Seymore Skinner

Entity

Contact.
Meet

Location

Entity

Springfield

Marge Simpson

per:cities_of_residence

per:children

per:children

per:spouse

BOTTOMLESS PETE, NATURE'S CRUELEST MISTAKE

per:alternate_names

Homer Simpson

Springfield Elementary

per:schools_attended

Lisa Simpson

Negative
Sentiment

Bart Simpson

TAC 2018 Planning Session: TAC 2018 Tracks

- Drug-Drug Interaction Extraction from Drug Labels
- Data Extraction for Systematic Review
- Streaming Multimedia Knowledge Base Population (joint with TRECVID)