



TAC KBP 2016 Linguistic Resources: Event Arguments (EA), Event Nuggets (EN) and Belief/Sentiment (BeSt)

Joe Ellis (presenter), Jennifer Tracey (presenter), Jeremy
Getman, Zhiyi Song, Ann Bies, Stephanie Strassel

*Linguistic Data Consortium
University of Pennsylvania*

- ◆ Linguistic resources for TAC KBP 2016
 - Eighth year LDC produced KBP resources
 - Twenty-nine new data sets
- ◆ Two primary goals
 - Increase coordination across tracks
 - Increase multi-lingual evaluation tracks
- ◆ Yesterday
 - Doc selection, ED&L, and Cold Start
- ◆ Today
 - Event Arguments, Event Nuggets, and Belief/Sentiment (BeSt)



- ◆ Entities, Relations, and Events (ERE)
 - Ongoing annotation task developed by LDC for DARPA's Deep Exploration and Filtering of Text program (DEFT)
 - Exhaustive labeling of entities, relations and events and their attributes.
- ◆ ERE annotation performed as upstream task
 - Provided inputs for multiple downstream tasks supporting ED&L, EA, EN, and BeSt
 - Primary means of meeting increased coordination of data goal

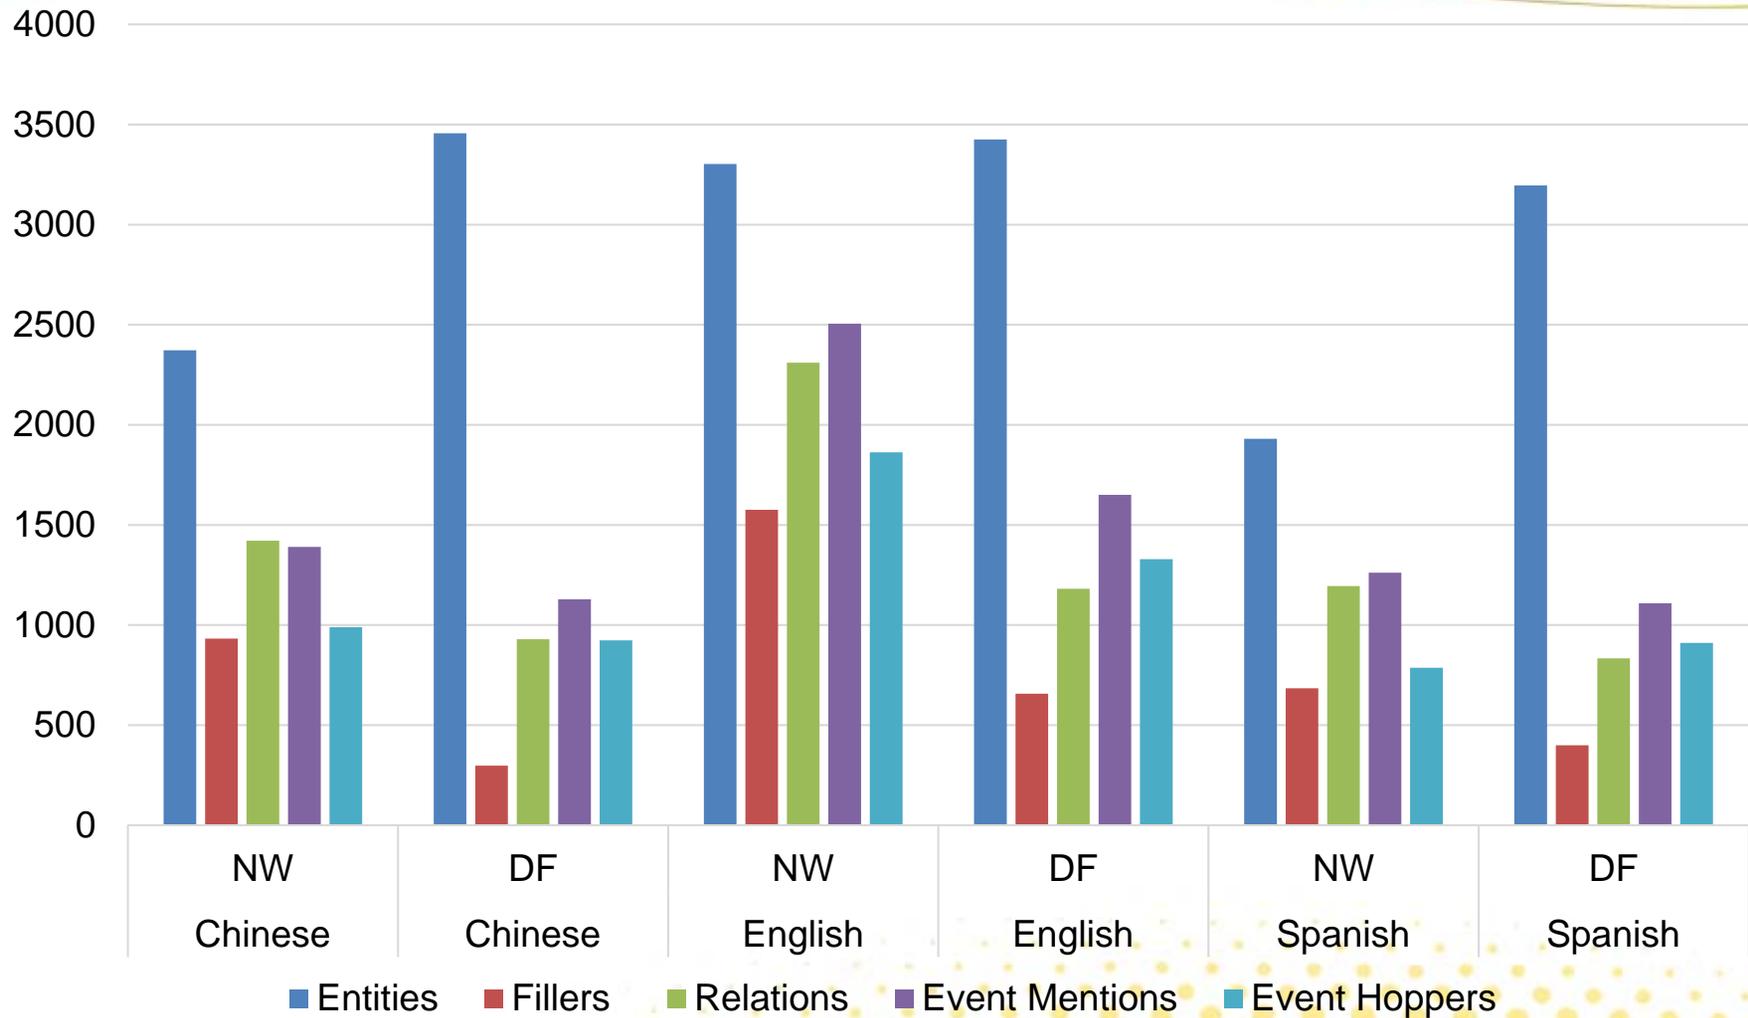
The **Bo Xilai** event was ignited by **Lijun running** into the **US consulate** in 2012 to bring **Bogu Kailai's killing** to light. Will **Bo Xilai end up** in **jail** due to bribery and corruption; what will **his wife** end up with?

Event				
	Type.Subtype	Realis	Trigger	Arguments
H1	Movement. Transport Person	Actual	running	Lijun
				US consulate
				2012
R2	Conflict. Personal.Family Attack	Actual	killing his wife	Bogu Kailai his wife
H3	Justice. Jail	Other	jail	Bo Xilai bribery and corruption
H4	Life.Die	Actual	killing	Bogu Kailai
H5	Movement. Transport Person	Other	end up	Bo Xilai
				jail

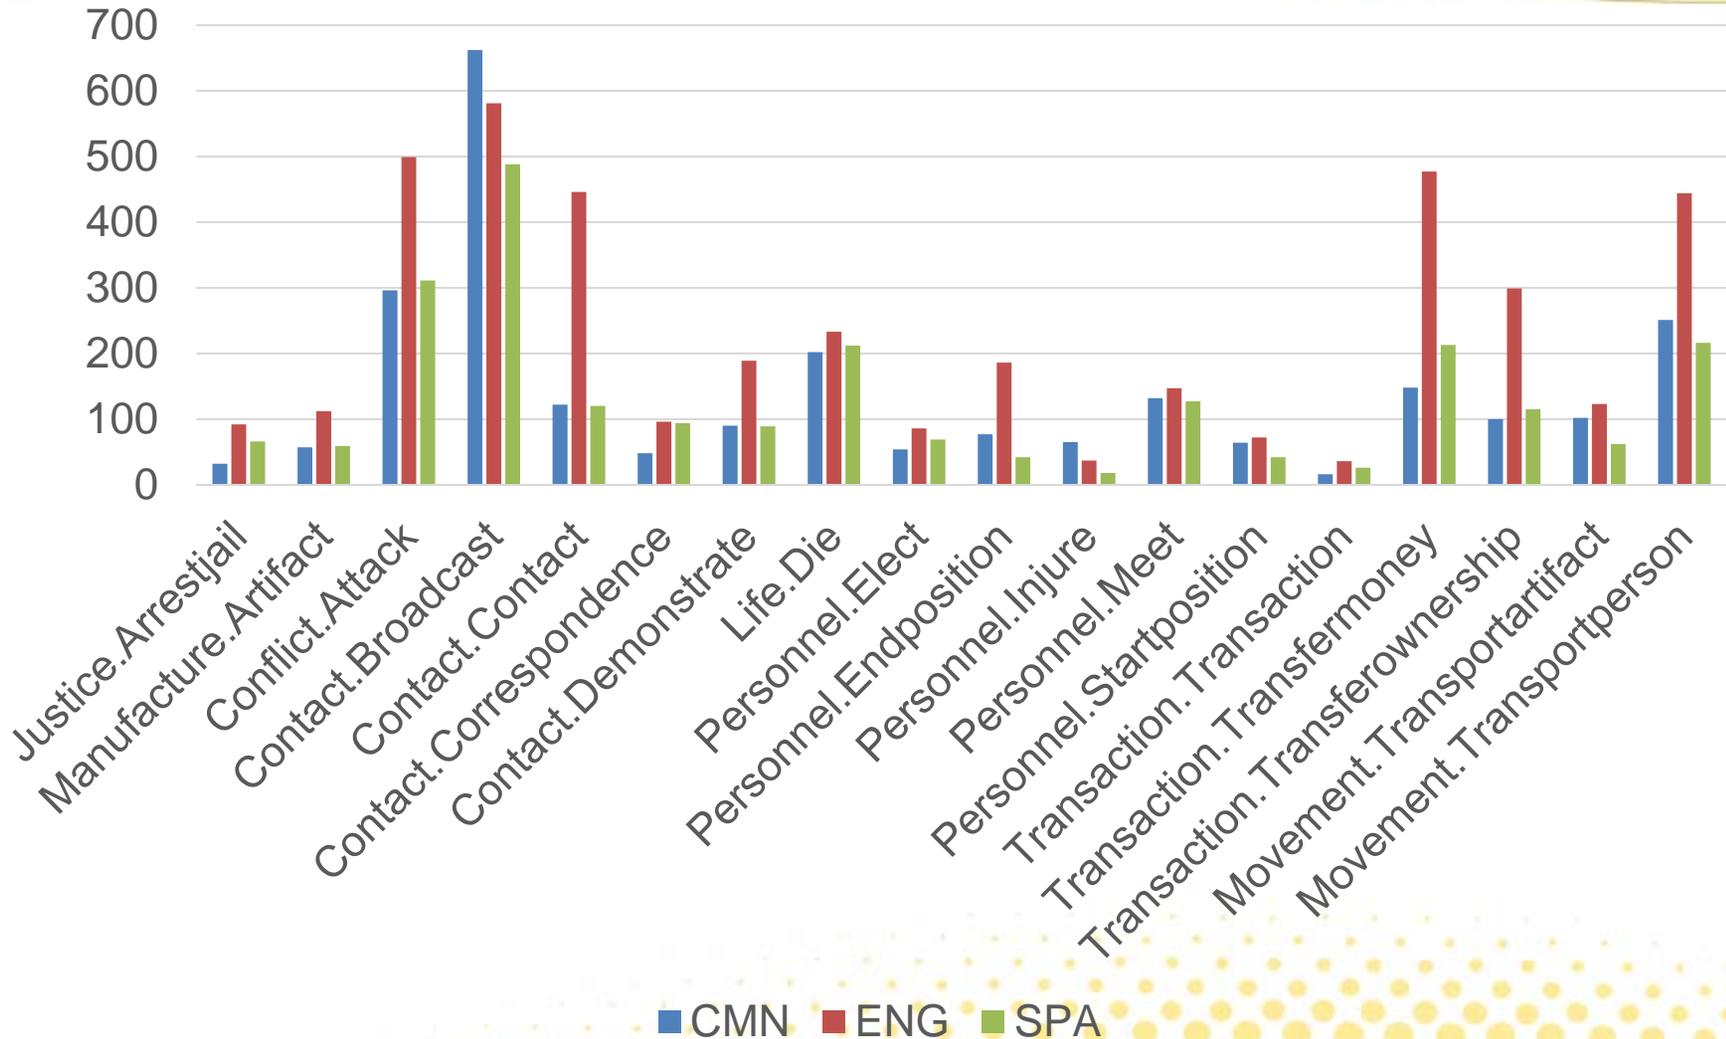
- ◆ ERE event type inventory reduced for 2016
 - 8 types and 18 subtypes in 2016 (listed below)
 - 9 types and 38 subtypes in 2015 (listed in overview paper)
 - Most of the dropped event types and subtypes are scarce in existing data

Conflict.Attack	Manufacture.Artifact	Justice.ArrestJail
Conflict.Demonstrate	Movement.TransportArtifact	Life.Die
Contact.Broadcast	Movement.TransportPerson	Life.Injure
Contact.Contact	Personnel.Elect	Transaction.Transaction
Contact.Correspondence	Personnel.EndPosition	Transaction.TransferMoney
Contact.Meet	Personnel.StartPosition	Transaction.TransferOwnership

ERE annotation counts



ERE event mentions



- ◆ Given new approach to evaluating EA in 2016, data development procedure overhauled
 - 2014-2015
 - Manual run
 - Argument-level assessment
 - 2016
 - Gold Standard
 - Event-level cross-document task
 - Queries, manual run, assessment

- ◆ ERE annotations on core source corpus
- ◆ Augmentation pass
 - BBN script run over ERE data
 - Annotators review results
 - Inferred arguments
 - Locational containment
 - Not annotated in ERE
 - Baghdad as Place of Conflict.Attack
 - Iraq added as 2nd Place

- ◆ Query selection
 - 51 simple, low-granularity queries
 - Event arguments in EAL Gold Standard
 - Annotators reviewed over 1300 potential queries
- ◆ Manual Run
 - Exhaustive across full 30K English source corpus
 - Justification strings indicating presence of event hopper in doc

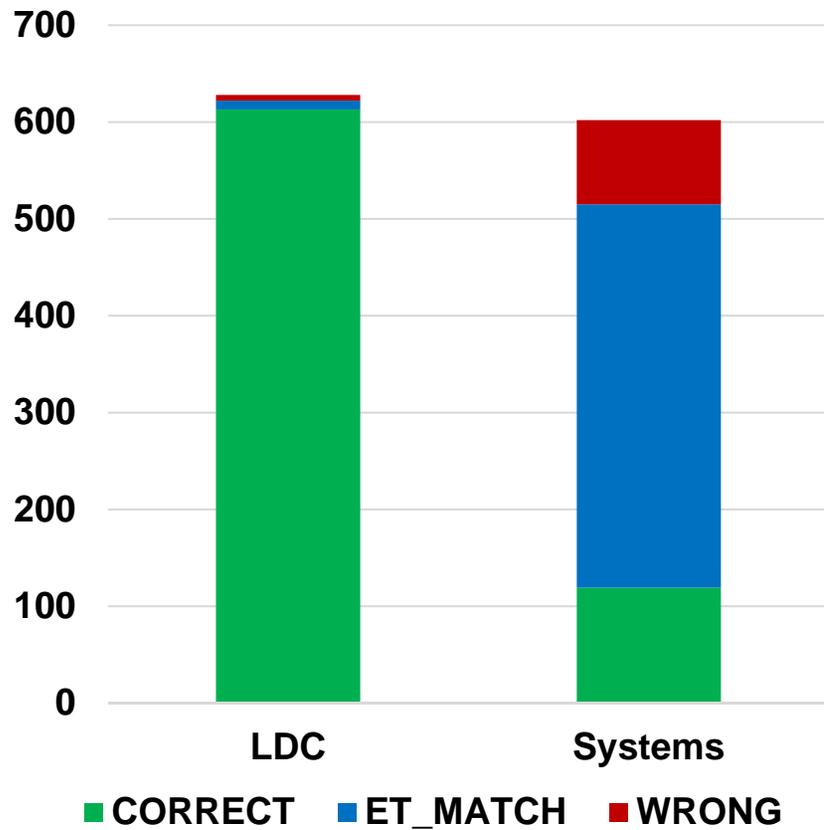
Earlier this week, former President Thabo Mbeki reportedly said South Africa needs better leadership...

Personnel.End-Position
Person – Thabo Mbeki

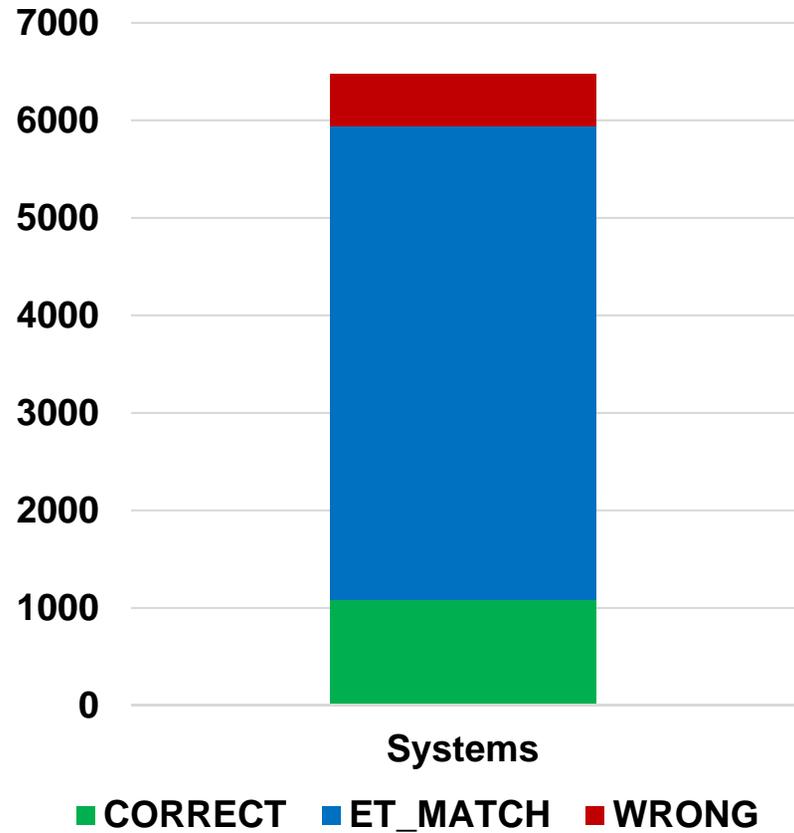
EA Cross-document Assessment and Results

- ◆ Does justification prove presence of query in document?
 - **Correct:** Response contains query event
 - **Event Type Match:** Contains event of same *type* as query, but not query event
 - **Wrong:** Doesn't contain query event or event of same type
- ◆ Low system recall on manually selected queries
 - BBN produced 249 “derived” queries based on system responses
 - No LDC manual run for these queries

EA Cross-Document Results



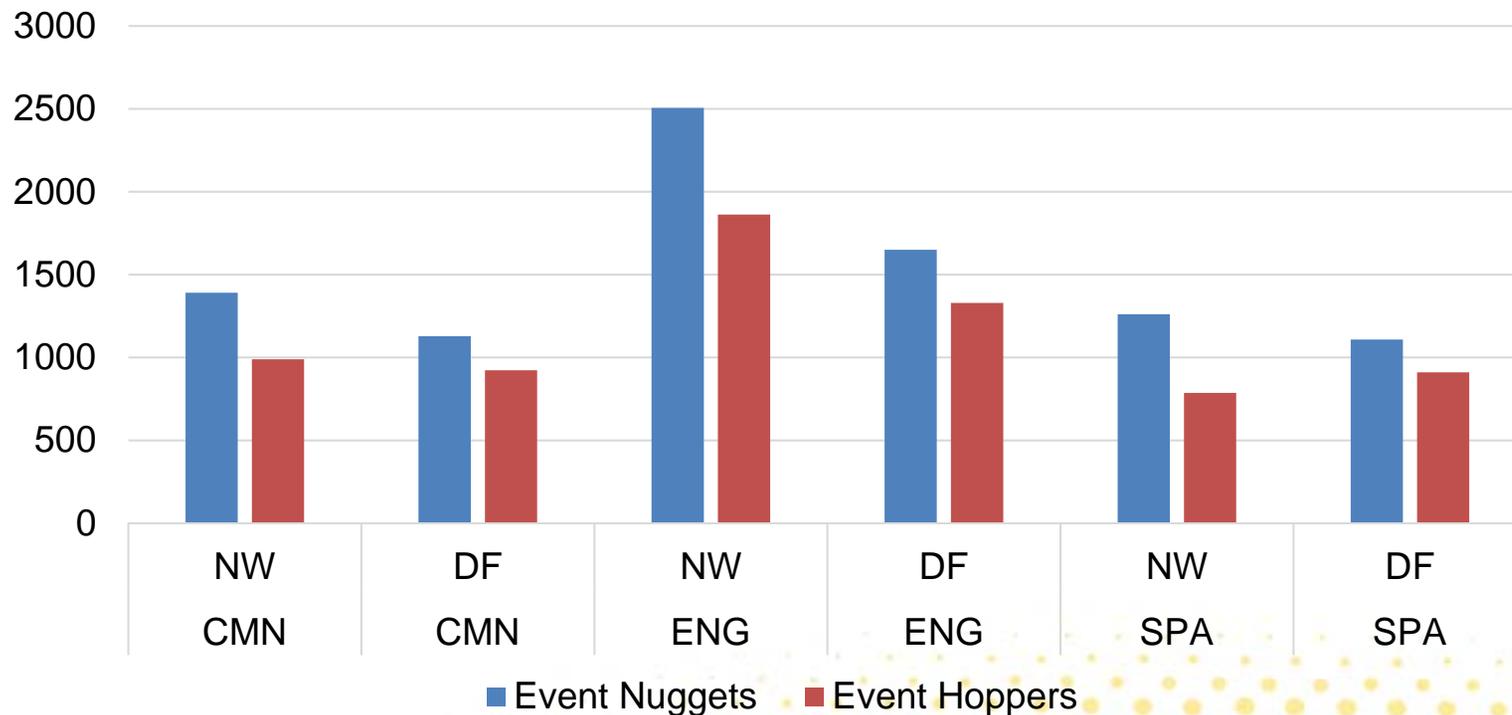
Manual Queries



Derived Queries

Event Nuggets and Linking (ENL)

- ◆ No separate ENL annotation task
 - Data are entirely produced by running a script over ERE data to extract and reformat a subset for use by ENL



Belief and Sentiment Annotation

- ◆ Only ERE entities are holders of belief and sentiment
- ◆ Only ERE entities, relations and events are targets of belief and sentiment
 - For events only, belief marked for each argument as well as the event itself
- ◆ Belief values: committed, non-committed, reported, n/a
 - Polarity also marked
- ◆ Sentiment values (polarity): positive, negative
- ◆ Sarcasm flag indicated when polarity annotated is opposite of literal meaning (based on context)

Belief Annotation Example

Ominous new action by **UKRAINE'S SECURITY FORCES** on Monday, including a **raid** on **AN OPPOSITION PARTY'S HEADQUARTERS**, appears to diminish prospects for **talks** between **PROTEST LEADERS**, a way to defuse the situation.

Event: raid

Event: talks

NA

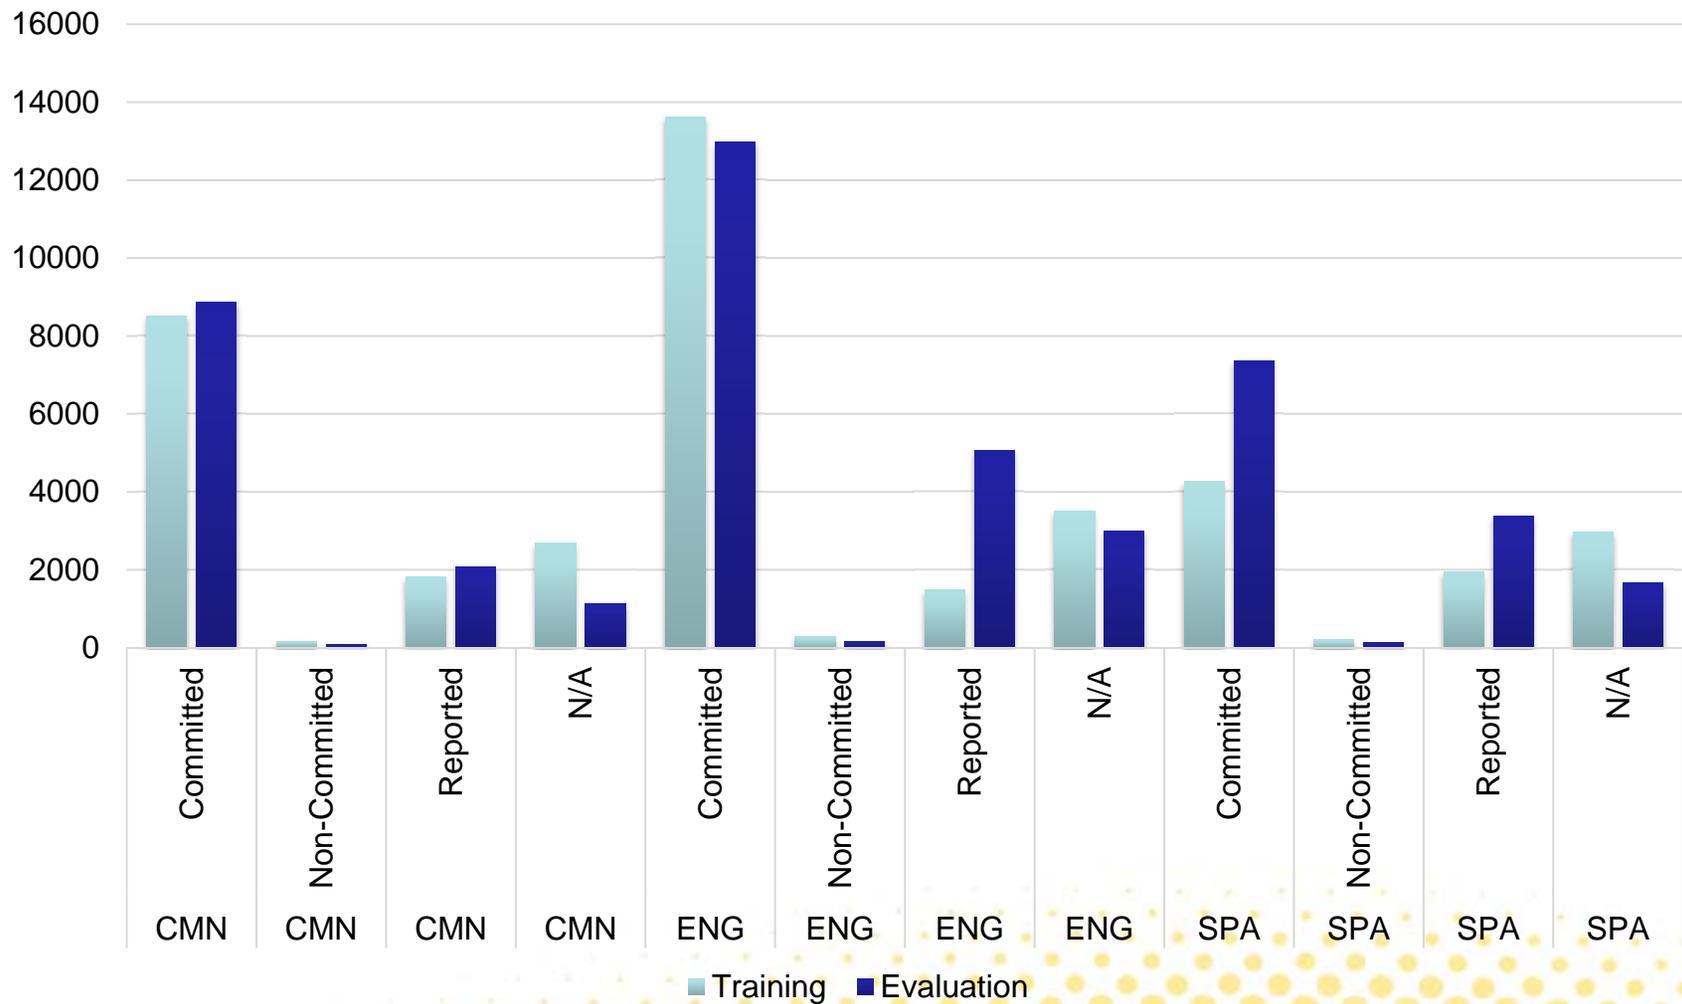
Arguments also

NA

Language	Belief annotations				Sentiment annotations			
	Training		Evaluation		Training		Evaluation	
	# Ann	#/Doc	# Ann	#/Doc	# Ann	#/Doc	# Ann	#/Doc
Chinese	13,192	66	12,163	76	27,982	140	18,982	118
English	18,915	77	21,188	128	38,664	157	25,358	154
Spanish	9,406	99	12,546	75	14,299	151	17,353	103

- ◆ English evaluation data notably more dense in belief annotations than training data
- ◆ Spanish evaluation data less dense than training data in both belief and sentiment

Comparison of Training and Eval Data - Belief



Comparison of Training and Eval Data - Sentiment

