LREC 2014 conference

Detecting Subevent Structure for Event Coreference Resolution

May 30, 2014

Jun Araki

Joint work with Zhengzhong Liu, Eduard Hovy, and Teruko Mitamura

Language Technologies Institute
School of Computer Science
Carnegie Mellon University

Outline

- Research problems with event coreference
- Subevent structure
- Our two-stage approach and results
- Error analysis
- Conclusion and future work

Research problems with event coreference

- Events can relate to each other in various ways
  - **Partial event coreference**; some event relations exhibit subtle deviation from perfect event identity (Hovy et al., 2013)

> In the town of Ercis, suspected rebels fired rockets at a police station, Anatolia said. No one was injured in the attack.

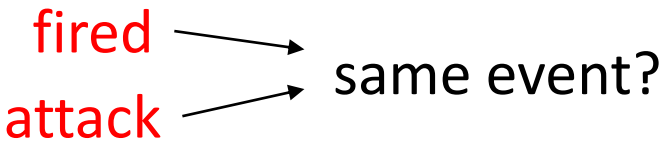

Subevents

- Mention 1 is a ***subevent*** of mention 2 if
  - mention 2 represents a stereotypical sequence of events, or a script, and
  - mention 1 is one of events executed as part of that script

Subevent structure

… when forces loyal to Egal's Ha-bar Awal sub-clan of the Issak attacked(E12) a militia stronghold of his main opposition rival, ...

Egal militia, claiming to be the national defence force, said they had captured(E15) two opposition posts, killing(E16) and wounding(E17) many of the fighters, destroying(E18) three technicals (armed pick-up trucks) and confiscating(E19) artillery guns and assorted ammunition.

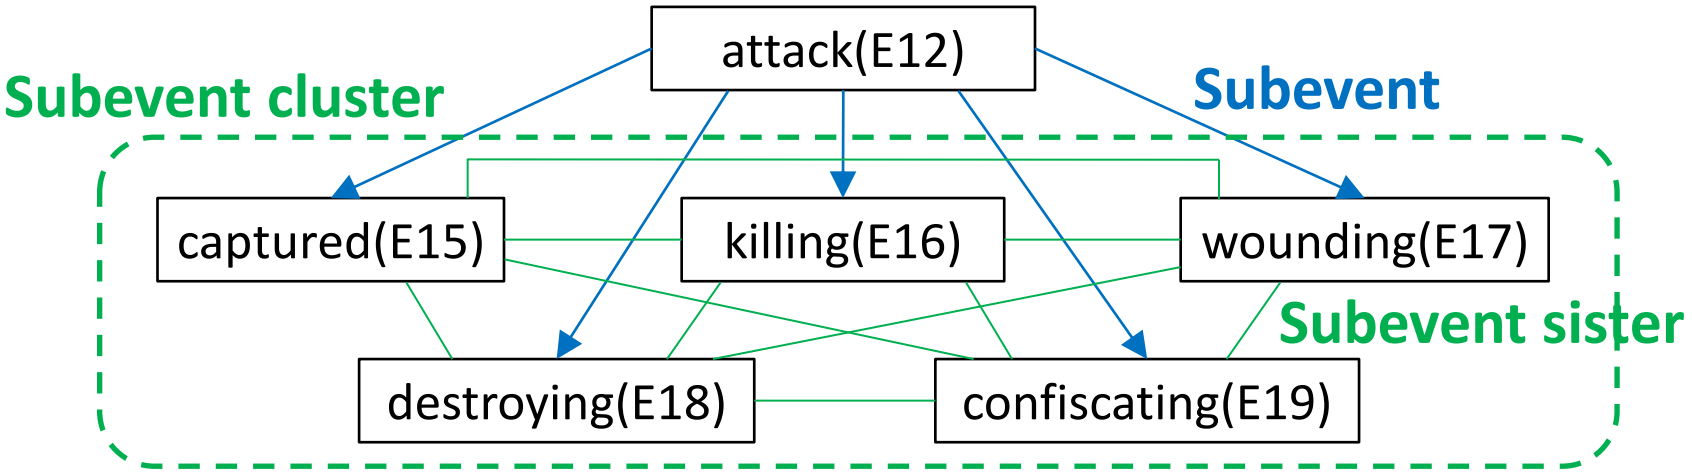

Corpus

- 65 newspaper articles in the violent domain
  - Event mentions are typically attacks, bombing, killing, etc.
- Inter-annotator agreement (Fleiss' kappa) (Hovy et al., 2013)
  - Full coreference: 0.620
  - Subevent: 0.467
- Propagated subevent relations are also counted as subevent relations

	Training+Dev	Test	Total
# Articles	49	16	65
# Relations	26499	9409	35908
FC	1037	216	1253
SP	997	201	1198
SS	399	139	538
NC	24066	8853	32919

Corpus statistics

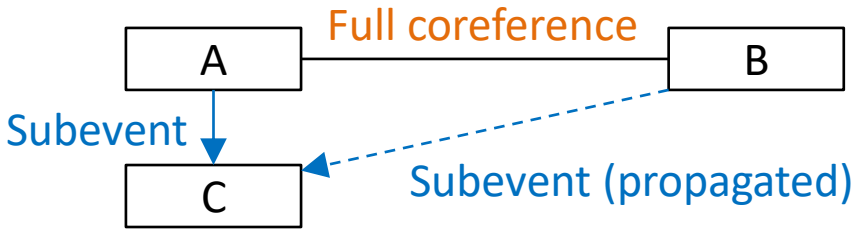

Two-stage approach

- Goal: Detecting subevent parent-child relations
- Our method
  - Basis: Pairwise coreference model (Chen et al., 2009; Bengtson and Roth, 2008)
  - Stage 1: Event relation learning
  - Stage 2: Subevent parent selection

Approach: stage 1

- Stage 1: Event relation learning
  - L2-regularized 4-class logistic regression model
  - Predicts one of the following classes for each pair
    - (1) Full coreference
    - (2) Subevent parent-child
    - (3) Subevent sister
    - (4) No coreference
  - 135 features from lexical, syntactic, semantic, and discourse levels
  - 5-fold cross validation

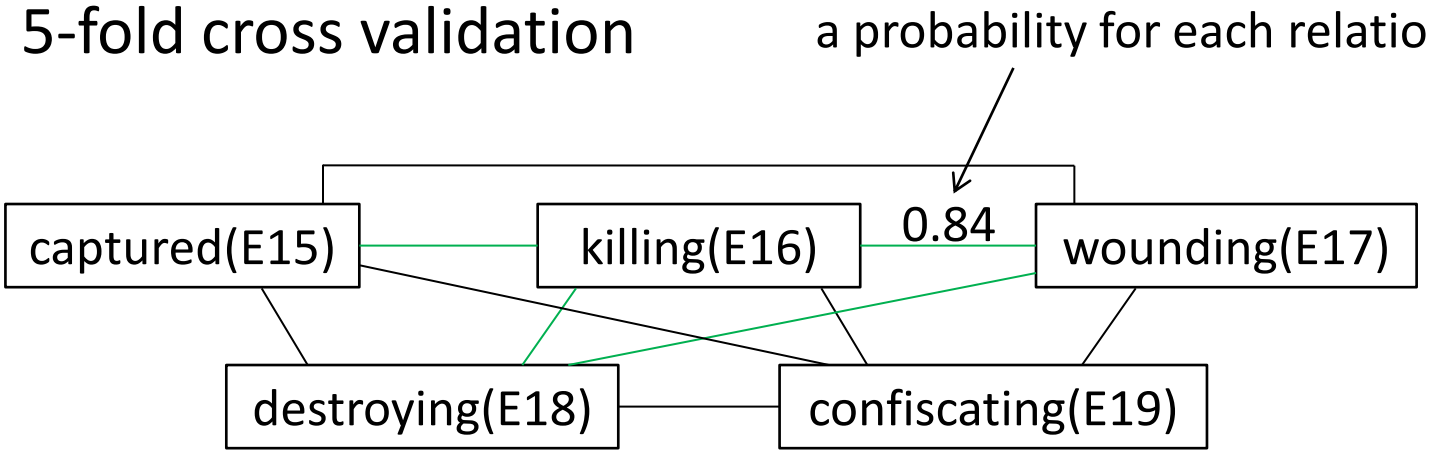

Experimental results: stage 1

- Evaluation
  - Apply BLANC (Recasens and Hovy, 2011) to 4 classes

$$F_{BLANC} = \frac{F_p + F_n}{2} = \frac{P_p R_p}{P_p + R_p} + \frac{P_n R_n}{P_n + R_n}$$

- Results
  - Difficult to gain high performance on subevent parent-child relations
  - Achieved high precision on subevent sister relations

Stage 1	BLANC				
	Pos links		Neg links		Avg
Relations	R_P	P_P	R_N	P_N	F1
Full coreference	41.20	41.59	98.64	98.62	70.01
Subevent parent-child	8.46	34.00	99.64	98.03	56.19
Subevent sister	14.39	66.67	99.89	98.73	61.49
No coreference	98.18	95.36	23.92	45.24	64.02

Approach: stage 2 (1)

- Stage 2: Subevent parent selection
  - Voting algorithm for selecting subevent parents
  - Option 1: $e_{sp} = \underset{e \notin sc, s \in sc}{\operatorname{argmax}} P_s(s, e)$

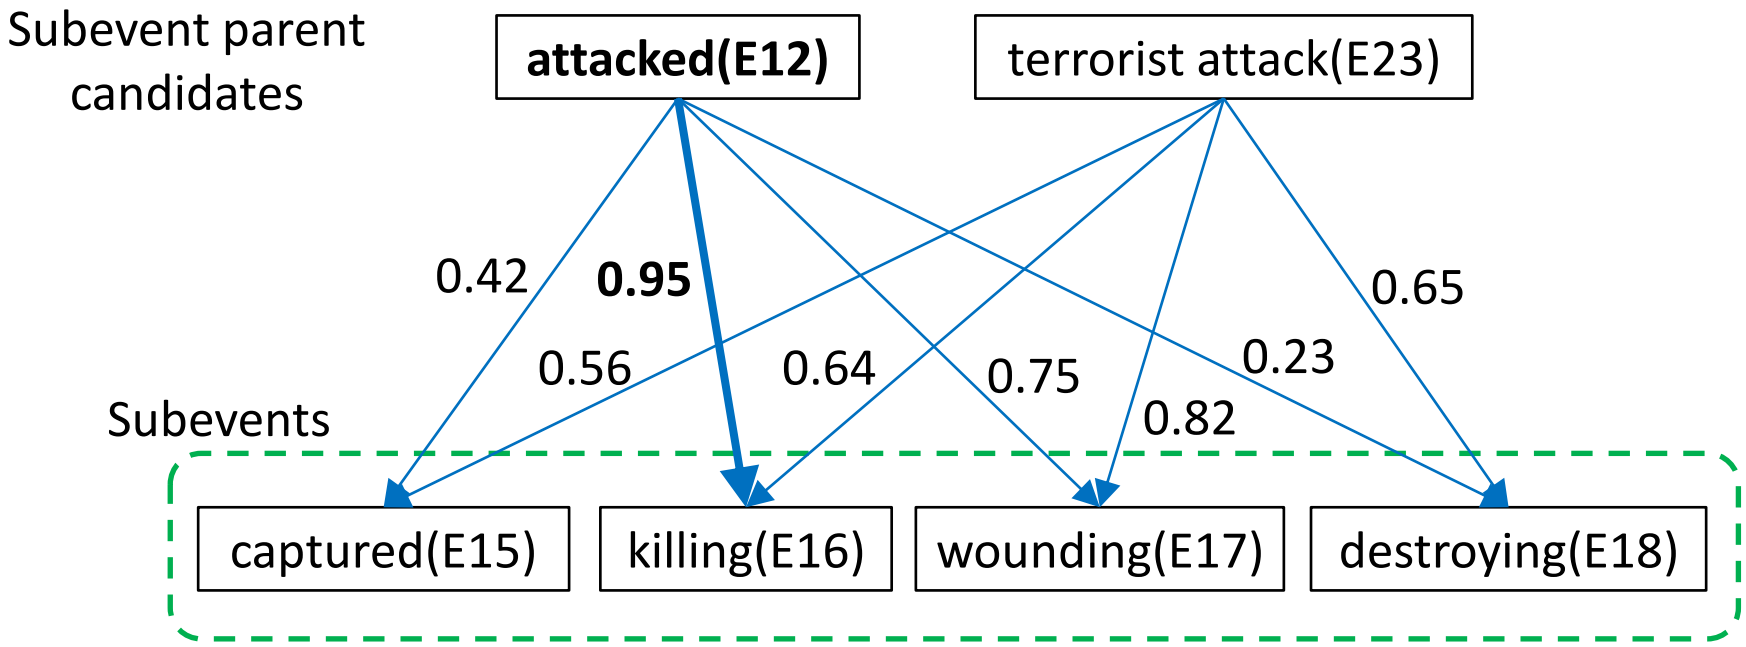

Approach: stage 2 (2)

- Stage 2: Subevent parent selection
  - Voting algorithm for selecting subevent parents
  - Option 2: $e_{sp} = \underset{e \notin sc}{\operatorname{argmax}} \sum_{s \in sc} P_s(s, e)$

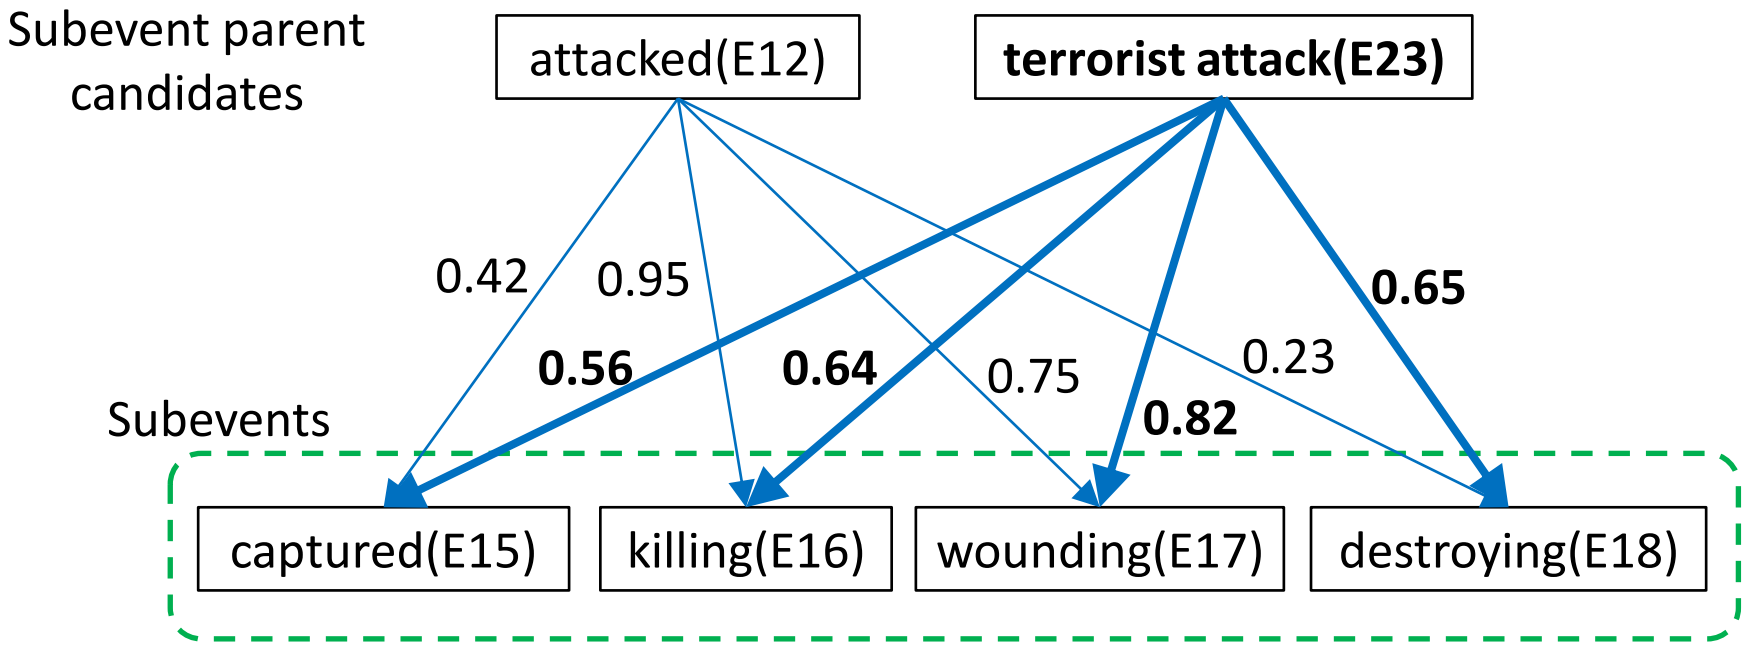

Experimental results: stage 2 (1)

- Stage 2 outperformed stage 1
  - Option 2 achieved better performance than option 1

	BLANC				
	Pos links		Neg links		Avg
Subevent parent-child	R_P	P_P	R_N	P_N	F1
Stage 1	8.46	34.00	99.64	98.03	56.19
Stage 2 (option 1)	13.43	31.03	99.35	98.13	58.74
Stage 2 (option 2)	14.43	33.33	99.37	98.15	59.45

Experimental results: stage 2 (2)

- Almost perfectly detected subevent structures

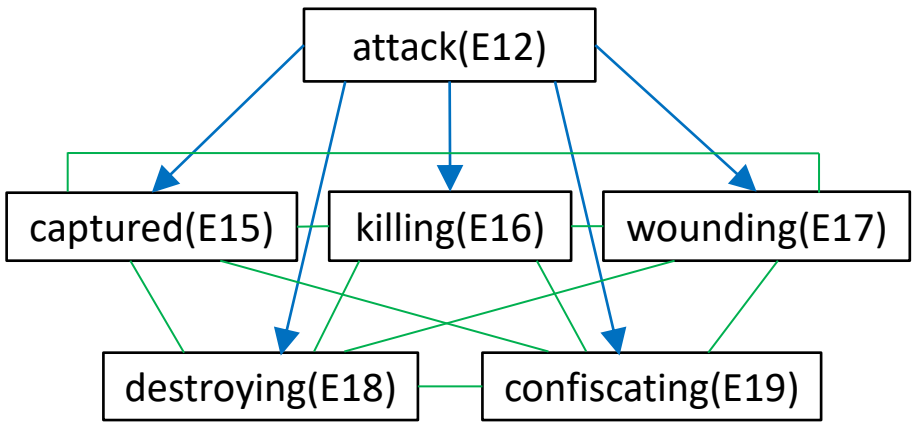

Gold standard

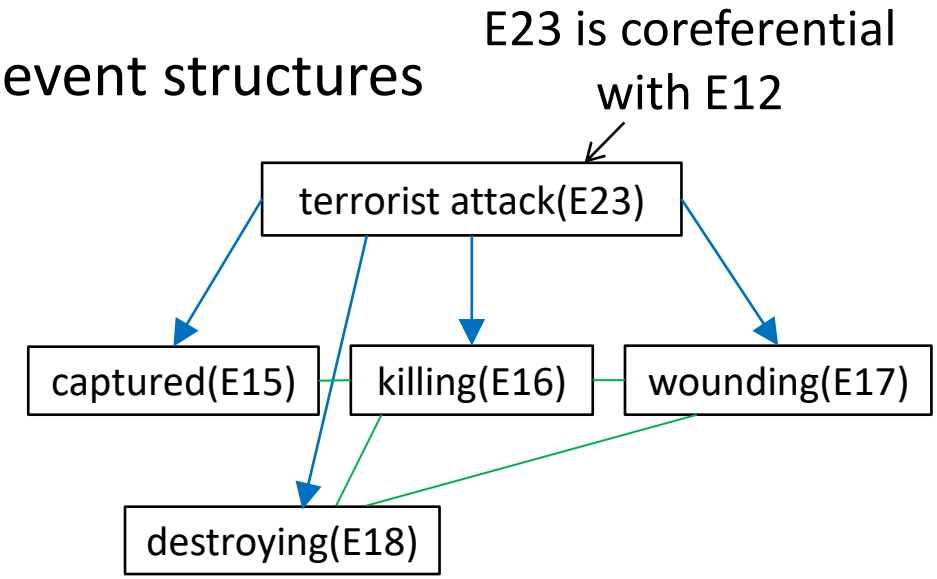

System output

- Perfectly detected subevent structures

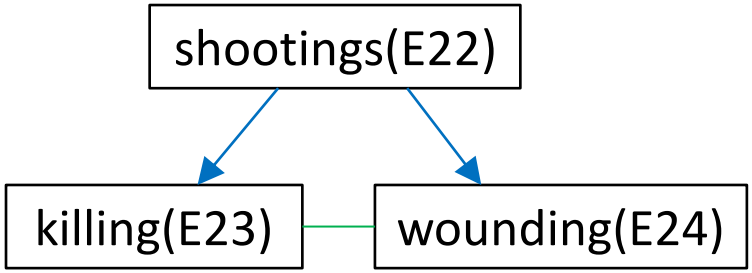

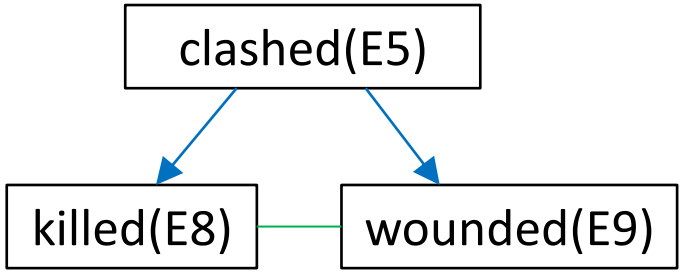

Error analysis

- Comparison between option 1 and 2
  - Some incorrect subevent parents gained a very high probability

> Violence(E21) also erupted in the West Bank, where Palestinian gunmen staged two shootings(E22), killing(E23) one Israeli man, Eldad Abir, 48, at a gas station, and seriously wounding(E24) a second man, the Israeli military said.

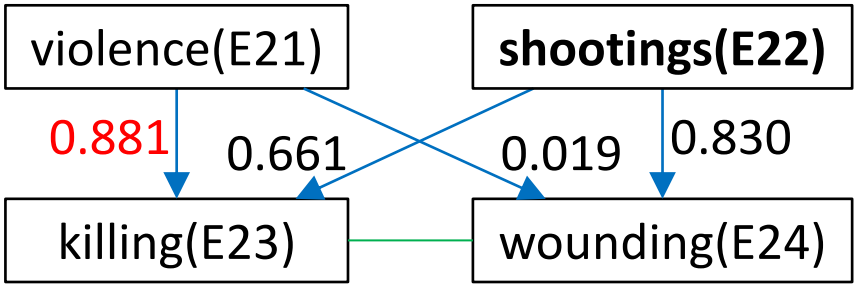

- Common errors
  - Linguistically complex expressions

> Over 90 Palestinians and one Israeli soldier have been killed(E14) since Israel launched(E15) a massive air and ground offensive(E16) into the Gaza Strip on June 28, ...

(E14 and E15 are subevents of E16)

Related work

- Most of event coreference work focuses on full event coreference
- Cybulska and Vossen (2012) considered partial coreference
  - Our work can capture subevent structure as well as subevents
- Some work (e.g., Chambers and Jurafsky (2008)) focuses on subevent sister relations, but not on subevent parent-child relations

Conclusion

- **Multi-class event coreference resolution**
  - Our logistic regression model can differentiate full coreference and subevent relations
    - It can also determine the directionality of subevent relations
- **Subevent structure detection**
  - We proposed a two-stage approach to improve subevent structure using a voting algorithm
    - It outperforms the logistic regression model on subevent detection

Future work

- Resolve structural inconsistency beyond pairwise decisions
- Deal with implicit subevent parents
  - They do not appear anywhere in text

> Six people were killed(E12) and 12 wounded(E13) when a suicide car bomber struck(E14) in Samarra, ...

(E12, E13, and E14 are subevents)

- Construct a library of domain event backbones

Thank you for your attention!