Carnegie Mellon University

Carnegie Mellon University Qatar

Interoperable Annotation of Events and Event Relations across Domains

Jun Araki, Lamana Mulaffer, Arun Pandian, Yukari Yamakawa, Kemal Oflazer, and Teruko Mitamura

Carnegie Mellon University

August 25, 2018 Interoperable Semantic Annotation Workshop @ Santa Fe, NM, USA

Motivation: Event structures

- Events are a core component for natural language understanding

A car bomb that police said was set by Shining Path guerrillas **ripped off**(E1) the front of a Lima police station before dawn Thursday, **wounding**(E2) 25 people. The **attack**(E3) marked the return to the spotlight of the feared Maoist group, recently overshadowed by a smaller rival band of rebels. The pre-dawn **bombing**(E4) **destroyed**(E5) part of the police station and a municipal office in Lima's industrial suburb of Ate-Vitarte, **wounding**(E6) 8 police officers, one seriously, Interior Minister Cesar Saucedo told reporters. The bomb **collapsed**(E7) the roof of a neighboring hospital, **injuring**(E8) 15, and **blew out**(E9) windows and doors in a public market, **wounding**(E10) two guards.

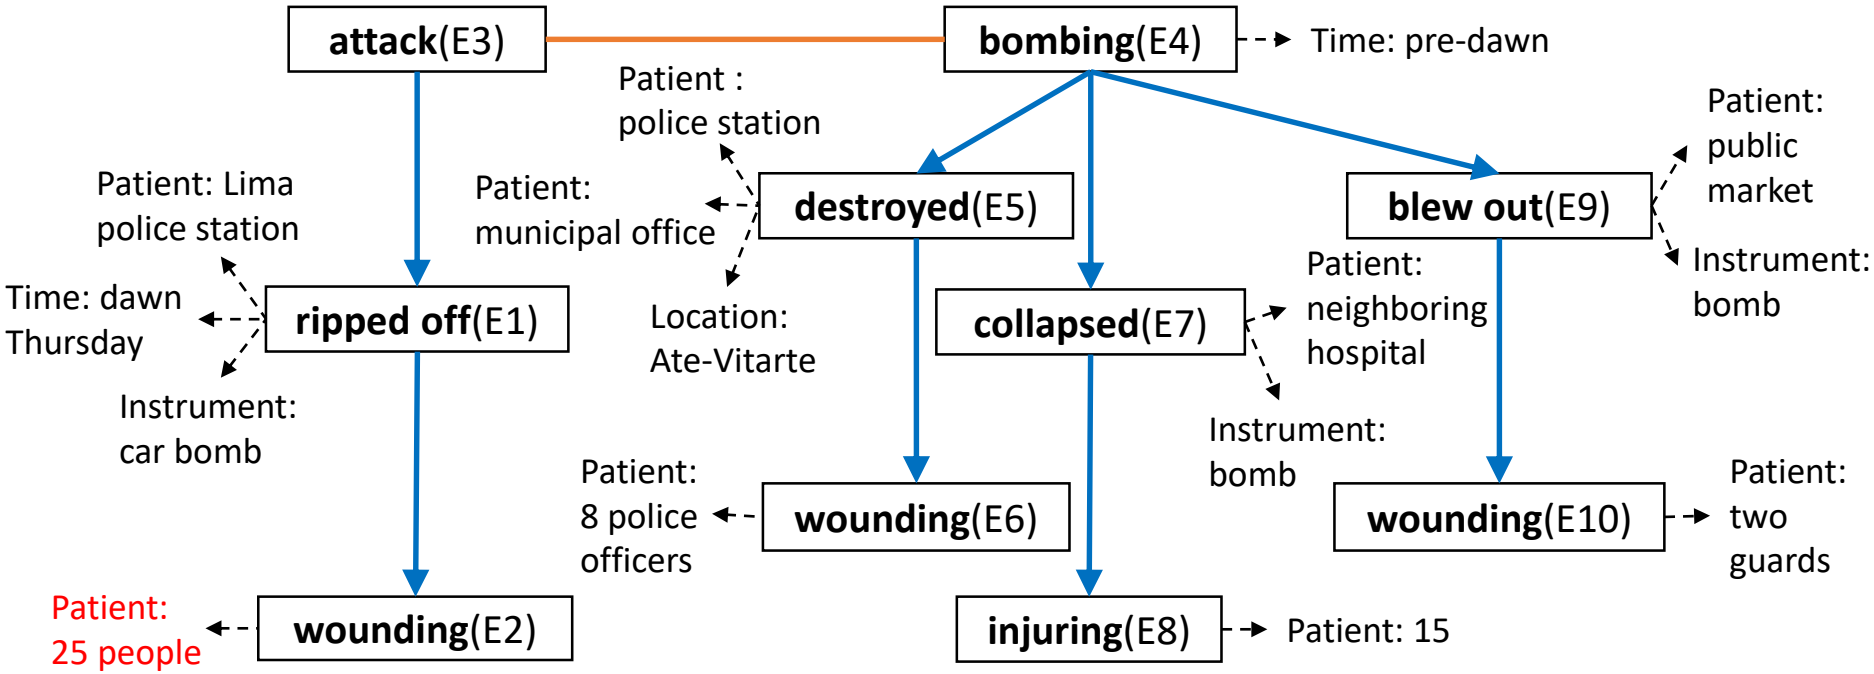

Motivation: Event structures

- Events are a core component for natural language understanding

The **attack**(E3) marked the return to the spotlight of the feared Maoist group, recently overshadowed by a smaller rival band of rebels. The pre-dawn **bombing**(E4) **destroyed**(E5) part of the police station and a municipal office in Lima's industrial suburb of Ate-Vitarte, **wounding**(E6) 8 police officers, one seriously, Interior Minister Cesar Saucedo told reporters. The bomb **collapsed**(E7) the roof of a neighboring hospital, **injuring**(E8) 15, and **blew out**(E9) windows and doors in a public market, **wounding**(E10) two guards.

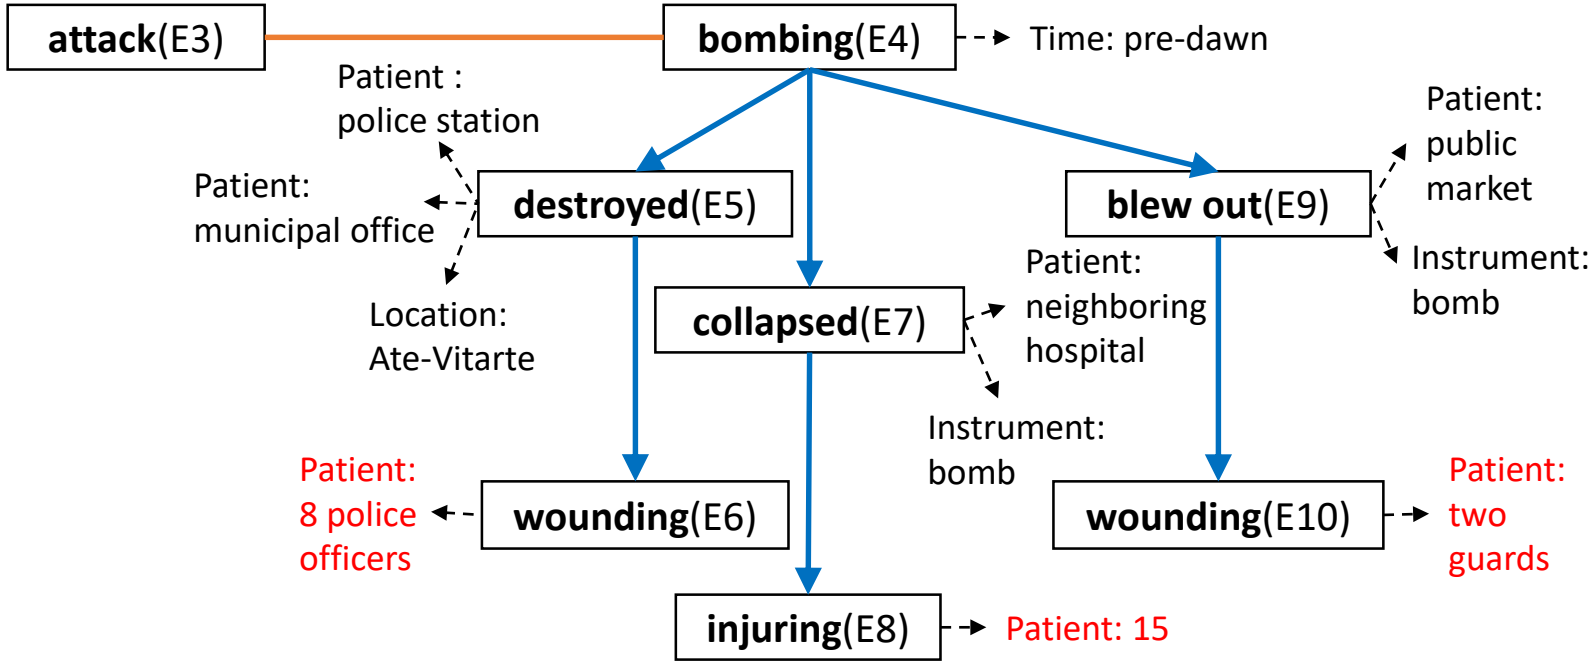

Event structures for question generation

- Generate high-level questions over multiple sentences via event relations
  - Require *inference steps* to resolve event relations
  - Useful to assess reading comprehension abilities of English-as-second-language (ESL) students [Araki+ 2016]
- Goal of this work
  - Provide human-annotated data to help us build question generation models

President Obama **met** with Putin last week.

Event coreference

The **meeting** took place in Paris.

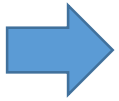

Q. Where did Obama meet Putin?

Jun Araki, Dheeraj Rajagopal, Sreecharan Sankaranarayanan, Susan Holm, Yukari Yamakawa, and Teruko Mitamura. Generating Questions and Multiple-Choice Answers using Semantic Analysis of Texts. COLING 2016.

Prior work on event annotation

- Closed-domain
  - Much work focuses on limited event types
    - MUC, ACE, TAC KBP, GENIA, BioNLP, and ProcessBank
- Open-domain
  - Some work focuses on conceptually different notions
    - WordNet, PropBank, NomBank, and FrameNet
  - Other work focuses on limited syntactic types
    - OntoNotes, TimeML, ECB+, and Richer Event Description (RED)

Our definition of events

- Eventualities [Bach 1986]
  - A broader notion of events
  - Consist of 3 components:

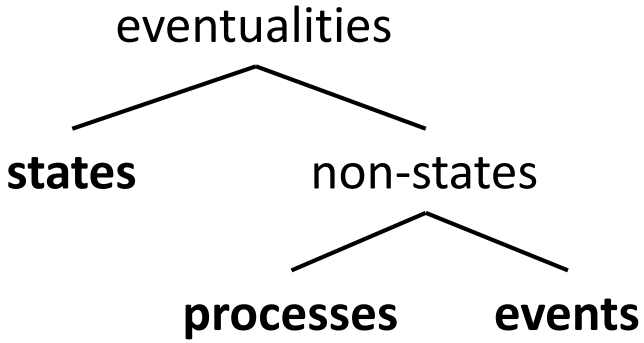

Component	Definition	Examples
states	a class of notions that are durative and changeless	want, own, love, resemble
processes	a class of notions that are durative and do not have any explicit goals	walking, sleeping, raining
events	a class of notions that have explicit goals or are momentaneous happenings	build, walk to Santa Fe, recognize, arrive, clap

Bach, E. The algebra of events. Linguistics and Philosophy, 9:5–16. 1986.

Our definition of events

- Event nuggets [Mitamura+ 2015]
  - A semantically meaningful unit that expresses an event
- Syntactic scope:
  - Verbs
    - Single-word verbs
    - Verb phrases
      - Continuous
      - Discontinuous
  - Nouns
    - Single-word nouns
    - Noun phrases
    - Proper nouns
  - Adjectives
  - Adverbs (+ verbs)

Examples:

The child **broke** a window …

She **picked up** a letter.

He **turned** the TV **on** … / She **sent** me an **email**.

The **discussion** was …

… maintained by **quality control** of …

Hurricane Katrina was …

She was **talkative** at the party.

She **replied dismissively** to …

Mitamura, T., Yamakawa, Y., Holm, S., Song, Z., Bies, A., Kulick, S., and Strassel, S. Event nugget annotation: Processes and issues. NAACL-HLT 2015 Workshop on Events: Definition, Detection, Coreference, and Representation.

Our definition of event relations

- **Event coreference**
  - A linguistic phenomenon that two event nuggets refer to the same event
  - Use the notion of *event hopper* from Rich ERE
- **Subevent**
  - Event A is a subevent of event B if B represents a stereotypical sequence of events, or a script [Schank+ 1977], and A is a part of that script

Examples:

The **Great Fire of London** happened in 1666. The **fire** lasted for three days.

New Orleans was **affected** by **Hurricane Katrina** which **flooded** most of the city when city levees **broke**.

Schank, R. and Abelson, R. 1977. Scripts, Plans, Goals, and Understanding: An Inquiry into Human Knowledge Structures. Lawrence Erlbaum Associates.

Our definition of event relations

- **Causality**
  - A cause-and-effect relation, in which we can explain the causation between two event nuggets X and Y, saying “X causes Y”
  - Inherently entails an event sequence
- Causality tests, based on [Dunietz+ 2017]
  - The “why” test
  - The temporal order test
  - The counterfactuality test
  - The ontological asymmetry test
  - The linguistic test
  - The granularity test

Examples:

The **tsunami** was caused by the **earthquake**.

Dunietz, J., Levin, L., and Carbonell, J. 2017. The BECauSE corpus 2.0: Annotating causality and overlapping relations. In Proceedings of the 11th Linguistic Annotation Workshop.

Our definition of event relations

- **Event sequence**
  - If event A is after event B, A happens after B happens under stereotypicality within a script or over multiple scripts
- **Simultaneity**
  - A relation that two event nuggets occur at the same time

Examples:

We **went** to **dinner** at a restaurant. We **ordered** steak and **ate** it. We then **got** a **call**. After the **call**, we **paid** and **left** the restaurant.

My boss was **talking** over the phone when I **stopped by** his office.

Overview of our annotation task

- **SW100**: Manually annotated 100 articles in Simple English Wikipedia
  - 10 different domains (e.g., geology and history)
  - 2 annotators and 1 more experienced annotator (adjudicator)
  - 5 event relations
    - event coreference, subevent, causality, event sequence, and simultaneity
- Steps:
  1. The 3 annotators identify event spans, following the annotation guidelines
  2. We compute inter-annotator agreement on event annotation
  3. The adjudicator finalizes event annotation.
  4. The 3 annotators identify event relations on top of the finalized events, following the annotation guidelines.
  5. We compute inter-annotator agreement on event relation annotation.
  6. The adjudicator finalizes event relation annotation.

Annotation tool: Our modified BRAT

- Original BRAT [Stenetorp+ 2012]
  - Stacks relation annotations vertically, which can deteriorate visualization significantly

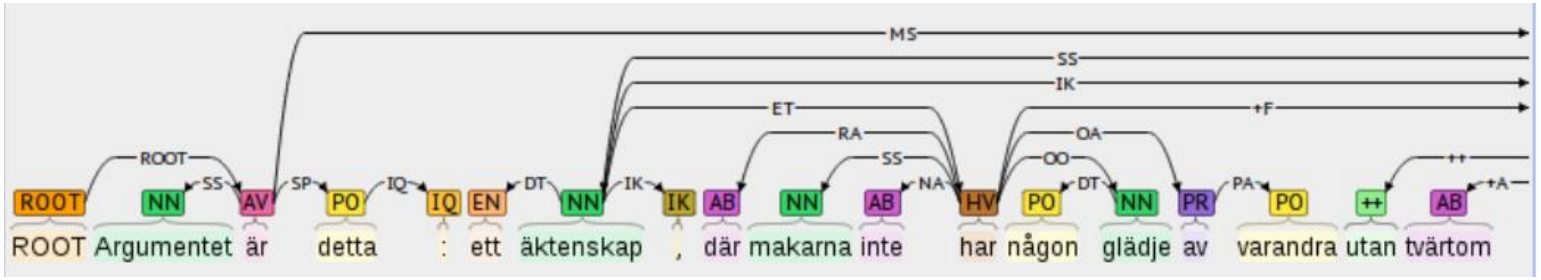

- Our modified BRAT
  - Improves visualization of relation annotations over multiple sentences

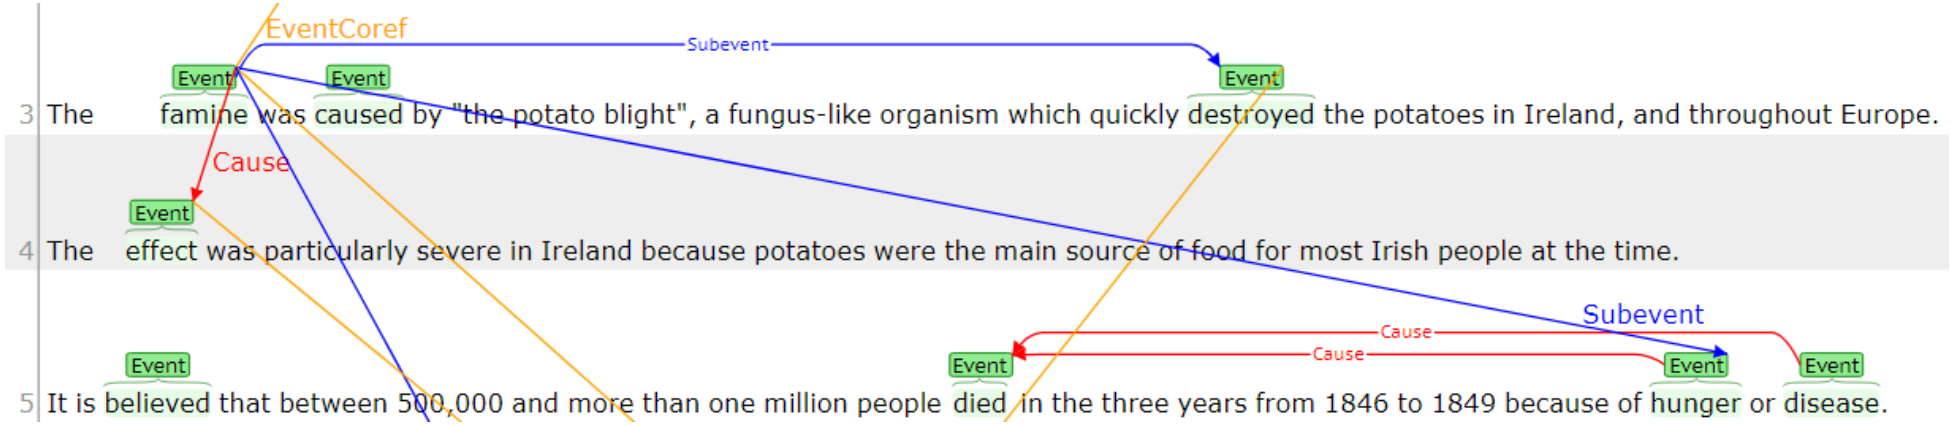

Stenetorp, P., Pyysalo S., Topic, G., Ohta, T., Ananiadou, S., and Tsujii, J. BRAT: A Web-based tool for NLP-assisted text annotation. EACL 2012: Demonstrations Session.

Corpus statistics of SW100

- Event annotation
  - 5,397 event nuggets

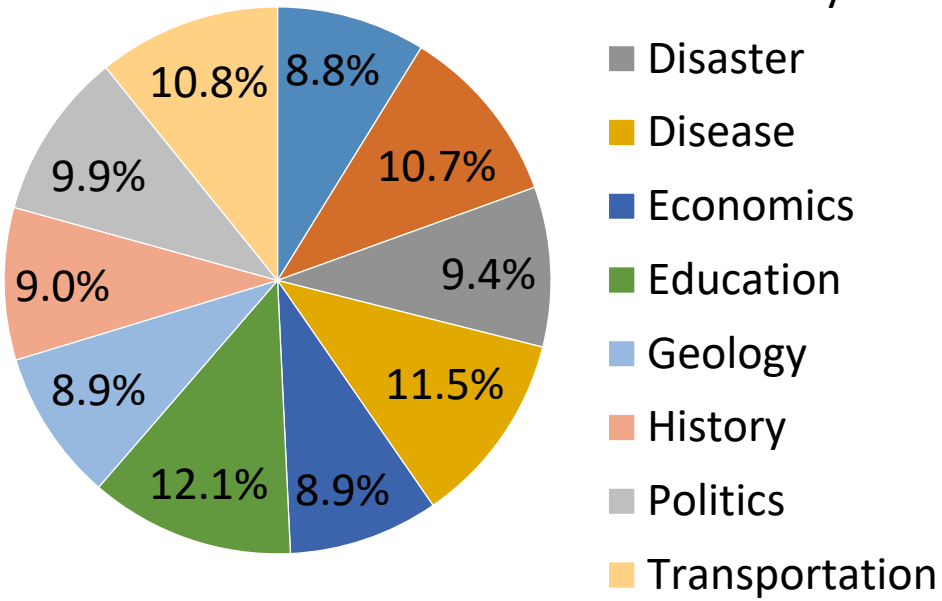

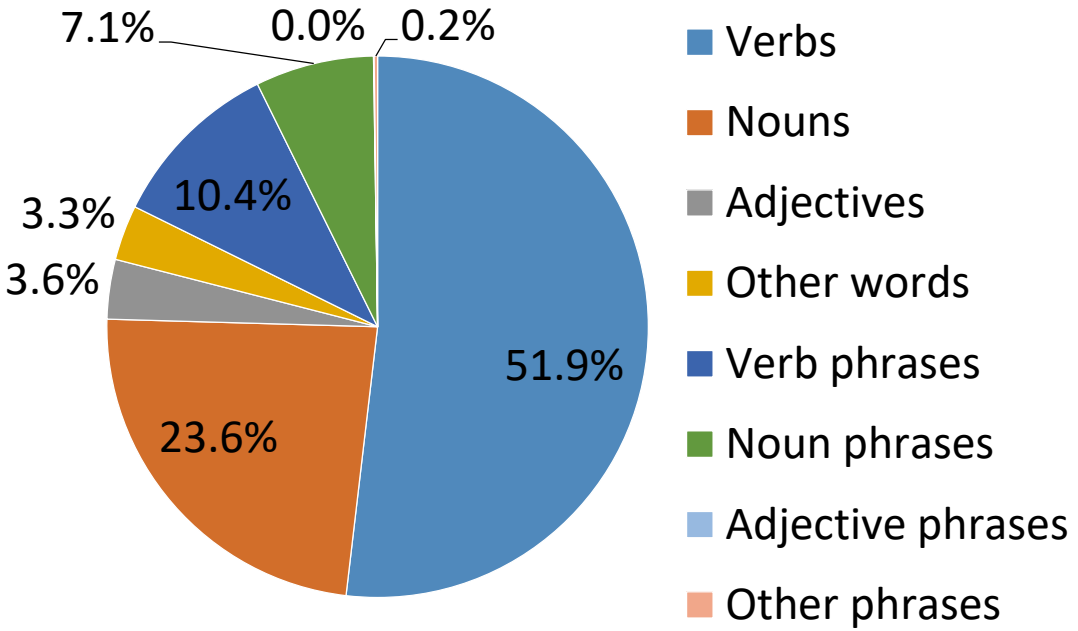

- Event relation annotation

	Arch	Chem	Disa	Dise	Eco	Edu	Geo	Hist	Poli	Tran	Total
# Event coreference clusters	32	61	49	67	51	61	52	42	46	51	512
# Subevent relations	74	81	98	94	48	101	91	119	147	93	946
# Causality relations	31	63	71	105	37	28	87	68	36	60	586
# Event sequence relations	48	38	63	30	73	62	58	117	66	88	643
# Simultaneity relations	8	11	21	10	15	17	8	21	24	11	146

Inter-annotator agreement on event annotation

- Measures inter-annotator agreement using the pairwise F1 score under two conditions
  1. Strict match: checking whether two annotations have exactly the same span
  2. Partial match: checking whether there is an overlap between annotations

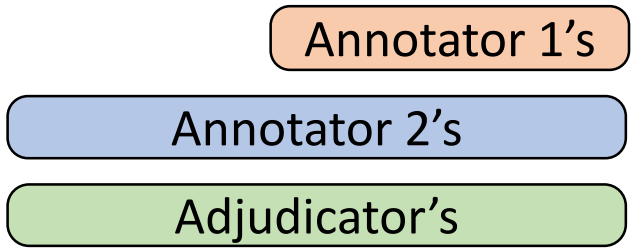

- Inter-annotator agreement = (average of two pairwise F1 scores)
  - **80.2%** (strict match) and **90.2%** (partial match)

Issues on annotation of events (1/2)

- Ambiguities on eventiveness
  - Examples:
    - These were **issues** of interest like the welfare state.
    - **Force** equals mass times acceleration.

  } "issues", "force" — an event or not?

  - We assume that there exists continuous semantic space between eventive and non-eventive

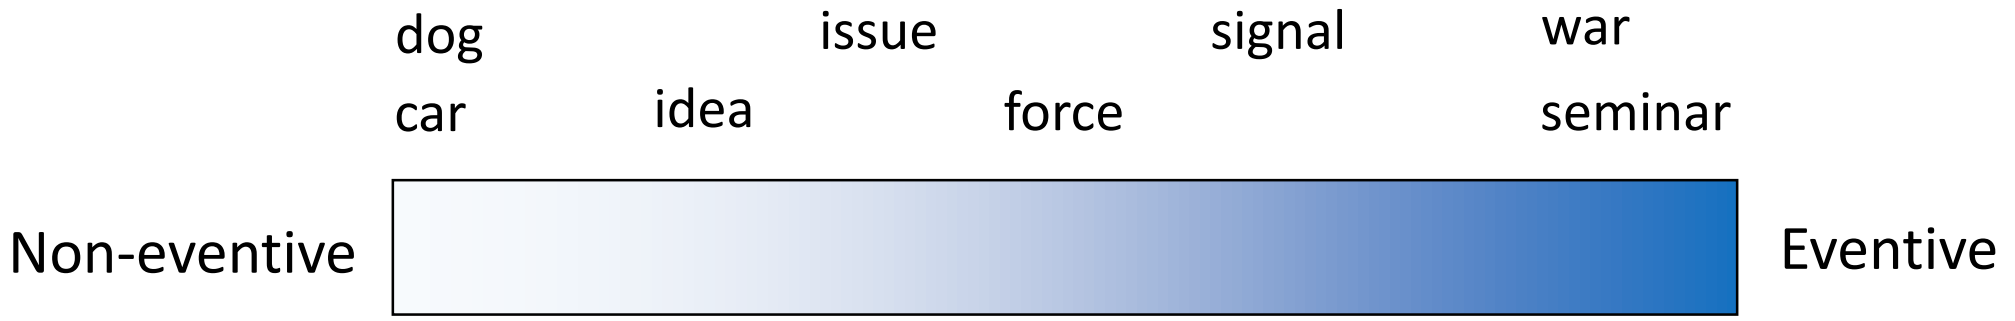

Issues on annotation of events (2/2)

- Ambiguities on semantic meaningfulness in the definition of event nuggets
  - Examples:
    - Bricks are used in masonry construction.

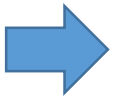

Is a “masonry construction” an event?

or

Is a “construction” an event?

a mere specifier (i.e., outside an event nugget)?

  - A clarification of “semantic meaningfulness” might be of some help

Inter-annotator agreement on event relation annotation

- Consider all pairwise relations between events and propagate event relations via event coreference
- Compute Fleiss' Kappa

Relation	*K*
Event coreference	0.645
Subevent	0.223
Causality	0.298
Event sequence	0.139
Simultaneity	0.108

Issues on annotation of event relations (1/3)

- Event annotation error
  - Error in event annotation can cause error in event relation annotation
  - "**Chronic Obstructive Pulmonary Disease** (**COPD**) can **make** breathing gradually **difficult**. **Breathing difficulties** caused by COPD can be compounded by ..."
    - If 'make breathing ... difficult' were annotated, it would be coreferent with 'Breathing difficulties'
- Event granularity
  - Needs to figure out a difference in event granularity between X and Y along with a certain script
  - "When Mount St. Helens **erupted** in 1980, it **released** 1000 times less material"
    - Two possible interpretations:
      1. 'Erupted' can be seen as a parent event of 'released' under the eruption script
      2. the two events have the same granularity and there is a causality relation

Issues on annotation of event relations (2/3)

- Script identification
  - The identification of scripts and their underlying subevents depends largely on common-sense knowledge and intuition of annotators
  - "He **sought treatment** for his **cancer**, after which he **got better**."
    - Is this a "sickness" script of (falling sick → recovering from it)?
- Domain-specific knowledge is required
  - "The 1973 **oil crisis** **started** on October 17, 1973, when the members of Organization of Arab Petroleum Exporting Countries (OAPEC) **said**, because of the **Yom Kippur War**, that they would no longer **ship** petroleum to nations that had **supported** Israel in its **conflict** with Syria and Egypt."
    - The correct annotation of event relations among these events requires comprehensive understanding of the 1973 oil crisis and the Yom Kippur War

Issues on annotation of event relations (3/3)

- Causality vs. Event sequence
  - We employ causality tests to differentiate causality from event sequence
  - “Igneous rock can **melt** into magma, **erode** into sediment, or be **pressed** tightly together to **become metamorphic**.”
    - A relation between ‘pressed’ and ‘become metamorphic’ could be causality or an event sequence
- Simultaneity vs. Event sequence
  - Our definition and annotation principles on simultaneity are not completely informative with respect to the duration of events
  - “When 1,500 missiles were **shipped**, three hostages were **released**.”
  - “A person can **have dyslexia** even if he or she is very smart or **educated**.”
    - The duration of events is often underspecified

Conclusion and future work

- Conclusion
  - Presented human annotation of events and five event relations of event coreference, subevent, causality, event sequence and simultaneity for the application of question generation in educational contexts
    - Formalized guidelines for annotating discontinuous event phrases
    - Achieved high inter-annotator agreement on event annotation, but lower inter-annotator agreement on event relation annotation
- Future work
  - Improve inter-annotator agreement on event and event relation annotation
    - Refine the definition of event nuggets and the differentiation between subevent relations, causality, event sequences, and simultaneity
    - Provide more comprehensive annotation guidelines would also lead to an improvement
  - Extend this work to languages other than English