This research was supported by Nuance Foundation Grant SC-14-74, NSF #IIS-1302668-002, NSF CISE CreativeIT #IIS-1002921 and NSF CISE RI EAGER #IIS-1044693.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the researchers and do not necessarily reflect the views of the National Science Foundation, Nuance Communications, or other sponsors of this project.

Inferring Narrative Causality between Event Pairs in Films

Zhichao Hu and Marilyn A. Walker
{zhu, maw}@soe.ucsc.edu
Natural Language and Dialog Systems Lab
University of California, Santa Cruz

Motivation & Background

- To understand narrative, humans draw inferences about the underlying relations between events
- Previous work either focused on "strict" physical causality [2][6], or event co-occurrence [1][3][4], and applied largely to newswire [1][3]
- We focus on Narrative Causality [7][8][9] - 4 types:
  - Physical Causality: Event A physically causes event B to happen
  - Motivational Causality: Event A happens with B as a motivation
  - Psychological Causality: Event A brings about emotions (expressed in event B)
  - Enabling Causality: Event A creates a state or condition for B to happen, A enables B

Frodo leaps to his feet and pushes his way towards the bar. Frodo **grabs** Pippin's sleeve, **spilling** his beer. Pippin **pushes** Frodo away...he **stumbles** backwards, and **falls** to the floor.

Bilbo leads Gandalf into Bag End... Cozy and cluttered with souvenirs of Bilbo's travels. Gandalf has to **stoop** to **avoid** hitting his head on the low ceiling. Bilbo hangs up Gandalf's hat on a peg and trots off down the hall. Bilbo disappears into the kitchen as Gandalf **looks** around.. **enjoying** the familiarity of Bag End... He **turns**, **knocking** his head on the light and then walking into the wooden beam.

Figure 1: Lord of the Rings, Fantasy Genre

Event Pair	Causality Type
grab - spill	Physical
push - stumble	Physical
push - fall	Physical
stoop - avoid	Motivational
look - enjoy	Psychological
turn - knock	Enabling

Table 1: Event pairs from Lord of the Rings scene with their causality types

Data & Method

- **Event**: non-stative verb with its arguments (generalized to person/something)
- **Event Pair:** two ordered events within a same document
- **955 film scene descriptions:**
  - 11 genres (# of films > 100 each genre)
  - 1.2 ~ 6.7 million words each genre
- Causal Potential [2]:

$$CP(e_1, e_2) = PMI(e_1, e_2) + \log \frac{P(e_1 \rightarrow e_2)}{P(e_2 \rightarrow e_1)}$$

$$\text{where } PMI(e_1, e_2) = \log \frac{P(e_1, e_2)}{P(e_1)P(e_2)}$$

  - Consists of two terms: pair-wise mutual information (PMI) and relative ordering of bigrams
  - Use a CPC (CP-Combined) measure
    - accounts for different window sizes
    - punishes event pairs from larger window sizes
    - w_{max}: max window size (3 in this paper); CP (e_1, e_2) using window size i

$$CPC(e_1, e_2) = \sum_{i=1}^{w_{max}} \frac{CP_i(e_1, e_2)}{i}$$

Experiments & Results

Download narrative causality event pairs!
https://nlds.soe.ucsc.edu/narrativecausality

High CPC Pairs	Low CPC Pairs
[person] clink [smth] - [person] drink [smth]	[person] strike - [person] give [person] [smth]
[person] beckon - [person] come	*[smth] become - [person] hide*
[person] bend - [person] pick up [smth]	*[person] lift [smth] - [person] cross*
[person] cough - [person] splutter	*[person] force - [smth] show [smth]*
High CPC Pairs	**Rel-gram Pairs**
[person] clear [smth] - [person] reveal [smth]	*[person] clear [smth] - [person] hit [smth]*
[person] embrace - [person] kiss	*[person] embrace [person] - [person] meet [person]*
[person] empty [something] -[person] reload	*[person] empty [smth] - [person] shoot [person]*
[person] stumble - [smth] fall	*[person] stumble upon [person] - [person] take [person]*

Table 2: High vs low CPC pairs from Exp 1, and high CPC vs top Rel-gram pairs from Exp 2, where high CPC pairs gained all Turkers' majority vote on a stronger causality relation

Experiment 1: High vs. Low CPC Event Pairs

- Top 3000 high CPC pairs from all genres - deduplicate to 960 pairs
- Bottom 6000 low CPC pairs from all genres
- Mechanical Turk: which is more likely to have a causal relation?
- Compare high-CPC pair to random low-CPC pair
- 5 Turkers, take majority vote
- Percentage of high-CP pairs labeled as causal: overall - 82.8%; Drama - 82.6%; Fantasy - **90.7%**; Mystery - **87.7%**; ……
- Smaller genres achieve higher causality rate

Experiment 2: CPC vs. Rel-gram [1] Event Pairs

- Rel-gram: *[police] arrest [person] - [person] be charge with [activity]*, with arguments generalized
- Mechanical Turk: which is more likely to have a causal relation?
- 100 random high pairs w. different first events
- Top Rel-gram pairs w. same first event as CPC pairs
- **81%** vote CPC pairs from film, 19% vote Rel-gram

Experiment 3: Narrative Causality Types

- Strong to weak: Physical, Motivational, Psychological, Enabling causality
- 100 random high-CPC pairs with all 5 Turkers' votes in Experiment 1
- Mechanical Turk: choose the strongest narrative causality
- 79% has majority vote: Motivational - 29%; Enabling - 28%; Physical - 13%; Psychological - 9%

Experiment 4: Genre Specific Causality

- 960 top pairs induced using separate genres vs. 960 top pairs induced using all films:
  - More than 70% overlap (with smaller genre sets most causal pairs were learned)
  - Mechanical Turk evaluation of non-overlap pairs shows quality of pairs from all vs. separate genres is similar
- 960 top pairs induced using separate ganres vs. 200 top pairs from camping & storm blogs [5]:
  - Only 2 overlap: *sit - eat, play - sing*
  - Causal relations learned from such small sets have topical and event-based coherence

References

- [1] Niranjan Balasubramanian, Stephen Soderland, Mausam, and Oren Etzioni. 2013. Generating coherent event schemas at scale. In Proceedings of EMNLP.
- [2] Brandon Beamer and Roxana Girju. 2009. Using a bigram event model to predict causal potential. In Computational Linguistics and Intelligent Text Processing, Springer.
- [3] Nathanael Chambers and Dan Jurafsky. 2008. Unsupervised learning of narrative event chains. Proceedings of ACL-08: HLT.
- [4] Karl Pichotta and Raymond J Mooney. 2014. Statistical script learning with multi-argument events. In EACL.
- [5] Elahe Rahimtoroghi, Ernesto Hernandez, and Marilyn A. Walker. 2016. Learning fine-grained knowledge about contingent relations between everyday events. In Proceedings of SIGDIAL.
- [6] Mehwish Riaz and Roxana Girju. 2010. Another look at causality: Discovering scenario-specific contingency relationships with no supervision. In ICSC.
- [7] Tom Trabasso, Paul Van den Broek, and So Young Suh. 1989. Logical necessity and transitivity of causal relations in stories. Discourse processes.
- [8] Paul Van den Broek. 1990. The causal inference maker: Towards a process model of inference generation in text comprehension. Comprehension processes in reading.
- [9] William H Warren, David W Nicholas, and Tom Trabasso. 1979. Event chains and inferences in understanding narratives. New directions in discourse processing.

Version: 08/08/2017