


Understanding Event Processes in Natural Language

Muhao Chen

Department of Computer Science / Information Sciences Institute

University of Southern California

Understanding Event Processes in Natural Language

How Do Machines Understand the *Evolution* of Events?

Human Language Always Communicates About Events

Earning a PhD in Computer Science typically takes around 5 years. It first involves ***fulfilling the course requirements*** and ***passing qualification exams***. Then within several years, the student is expected to ***find a thesis topic***, ***publish several papers*** about the topic and ***present them in conferences***. The last one or two years are often about ***completing the dissertation proposal***, ***writing*** and ***defending the dissertation***.

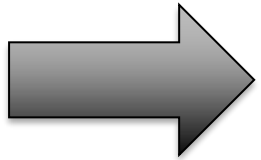

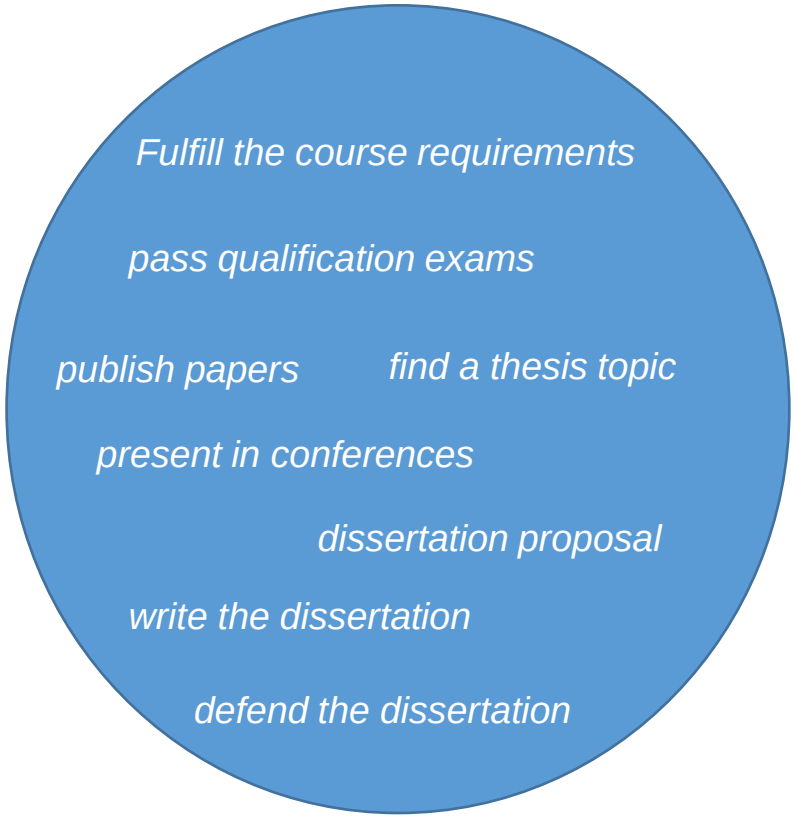

Natural language understanding (NLU) has to deal with event understanding

Event Extraction

What is an event?

An action or **a series of actions** that happen at a specific location, within **a period of time**, and causes change(s) to the status of **some object(s)**

E.g.:

Jeff *shaved* my hair yesterday at home

How to recognize an event in text?

Supervised Methods

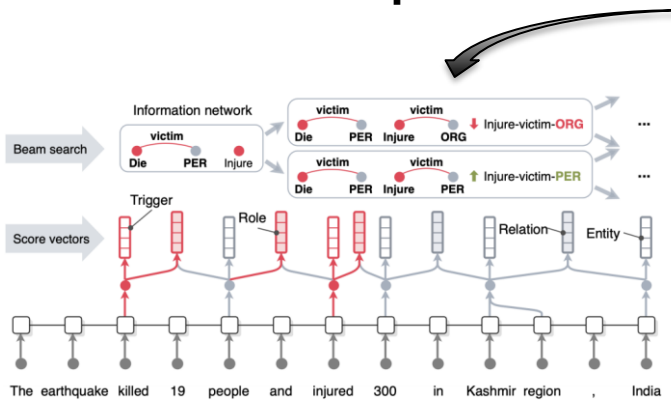

Bi-LSTM-CRF, Seq2Struct, etc.

Annotated documents

- E.g, ACE-05, RED, ERE, etc…

Unsupervised Methods

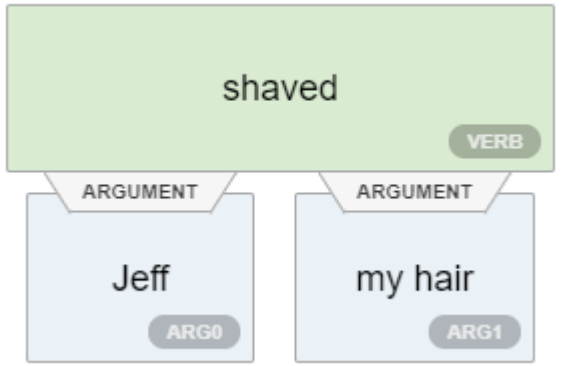

Semantic Role Labeling
(Verb SRL / Nom SRL)

Event Process Understanding And Prediction

Extraction only is not enough.

Events are **NOT simple, static** predicates.

They evolve,

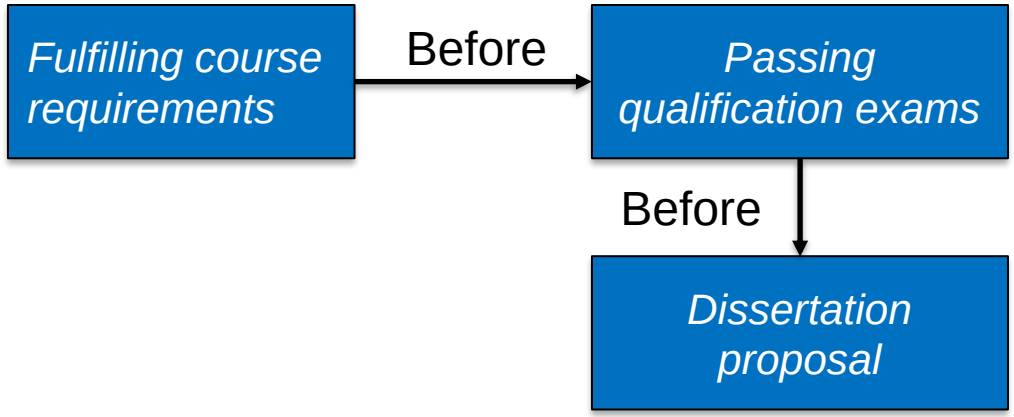

are described in different granularities,

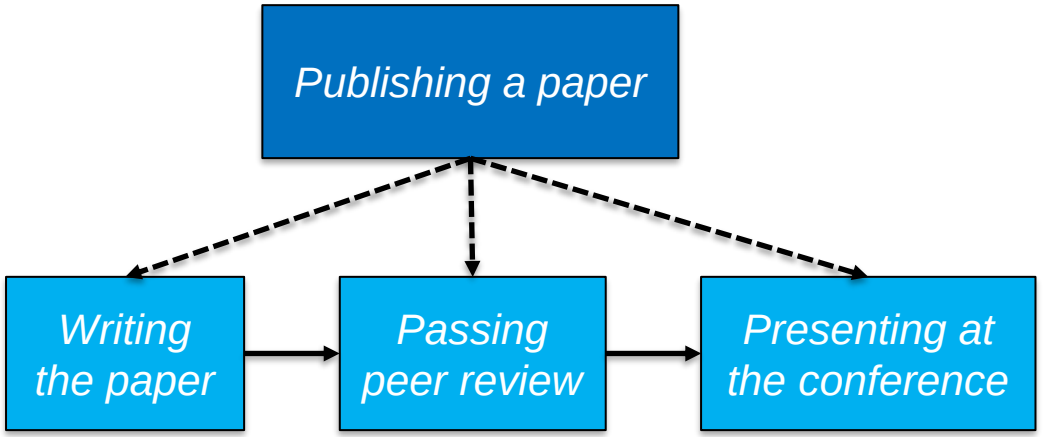

and are always directed by specific intents or central goals [Zacks et al. *Nature Neuroscience*, 2001]

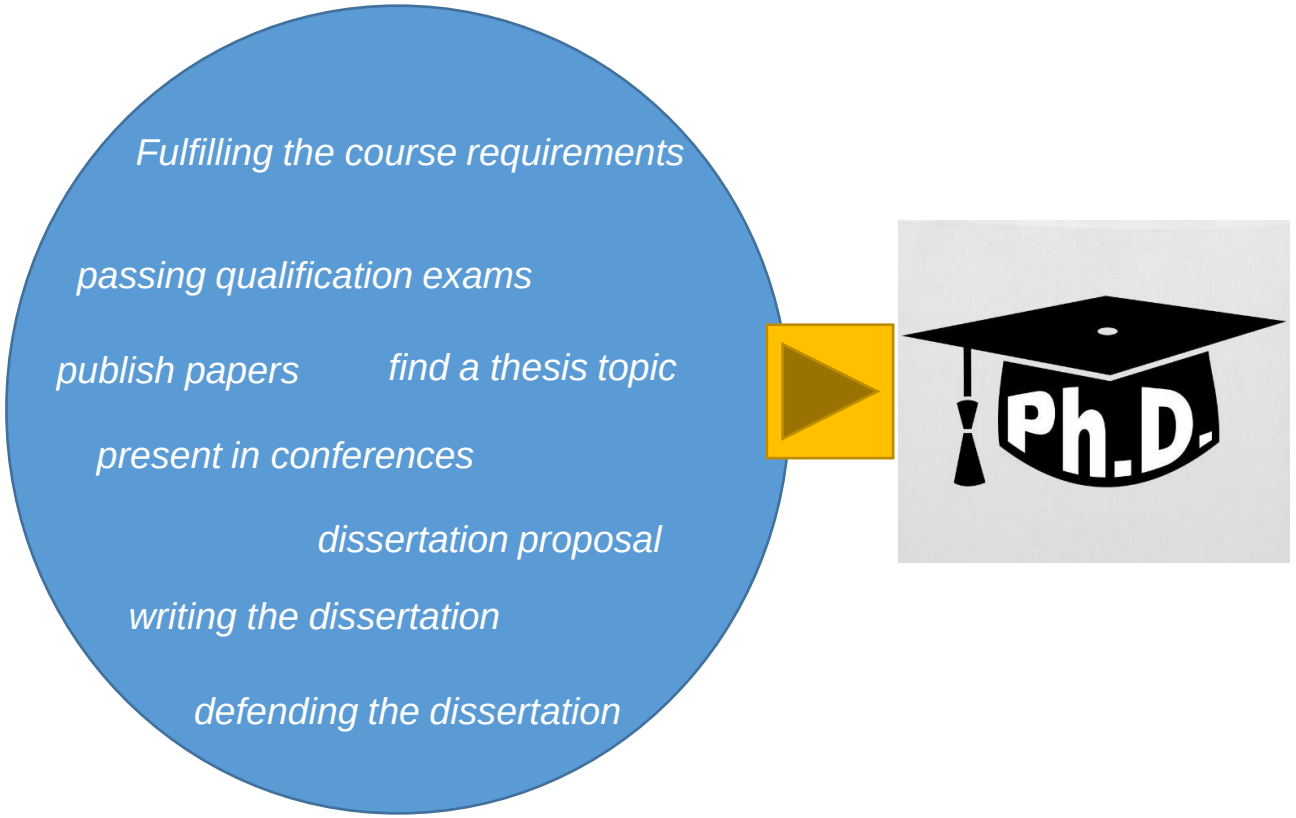

Event Process Understanding And Prediction

An event process (or event chain)

- Partially ordered events that are centered around common protagonists [Chambers et al., ACL-08]

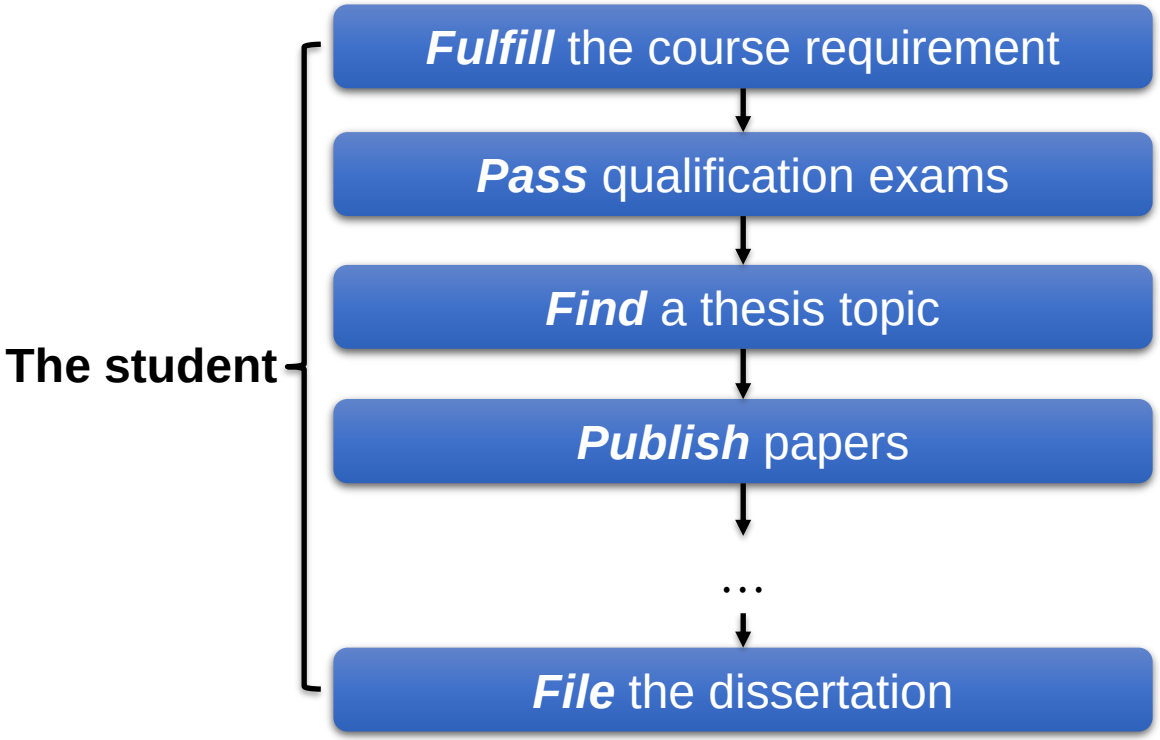

Prediction problems on event processes

Event process completion

- What happens next?

Intention prediction

- What is the goal of "*digging a hole, putting some seeds in the hole* and *filling it with soil*"?

Membership prediction

- What are the steps of "buying a car"?

Salience prediction

- Is *defending the dissertation* more important than *doing an internship*?

Event Processes Are Essential to Downstream NLU Tasks

Narrative prediction

One day Wesley's auntie came over to visit. He was happy to see her, because he liked to play with her. When she started to give his little sister attention, he got **jealous**. He got **angry** at his auntie and **bit** his sister's hand when she wasn't looking.

Then what might happen?

O1: He was **scolded**. ✓

O2: She **gave him a cookie** for being so nice. ✗

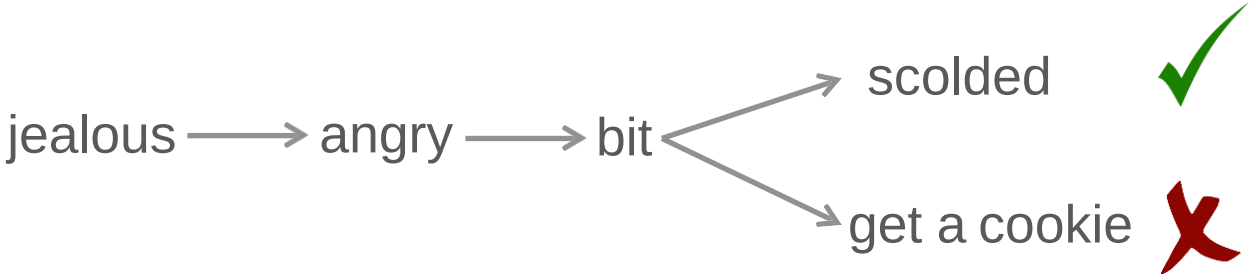

Machine comprehension

Water is split, providing a source of electrons and protons (hydrogen ions, H^+) and giving off O_2 as a by-product. ***Light absorbed*** by chlorophyll drives a ***transfer of the electrons and hydrogen ions*** from water to an acceptor called $NADP^+$.

What can the splitting of water lead to?

A: Light absorption

B: Transfer of ions

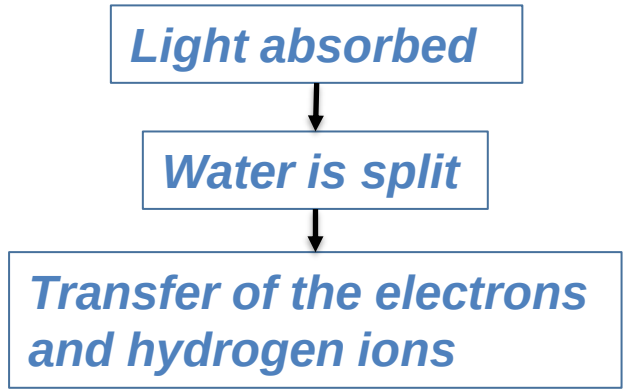

Agenda

1. Event process completion

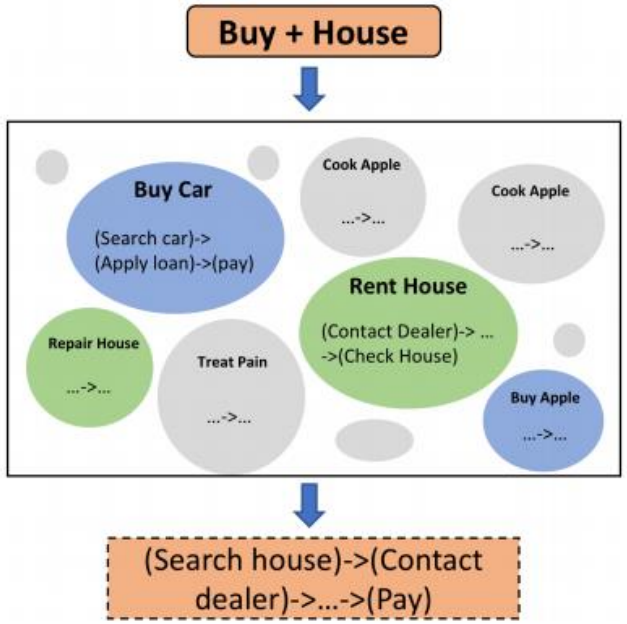

2. Event intention prediction

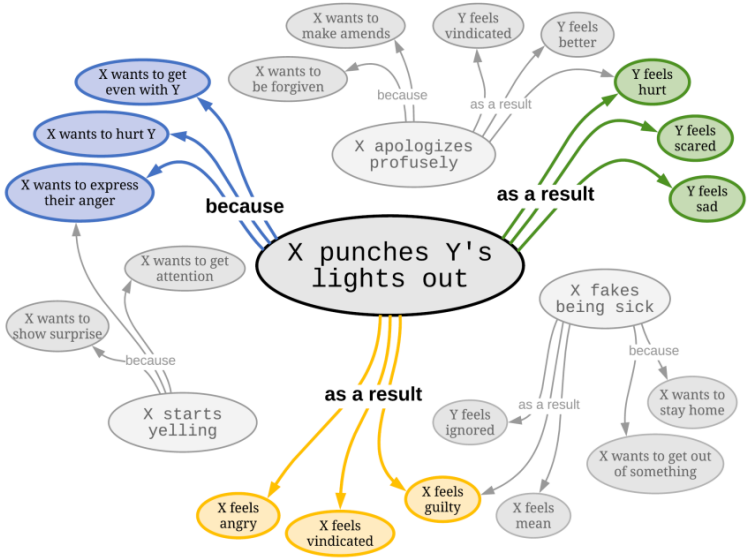

3. Event processes in downstream NLU tasks

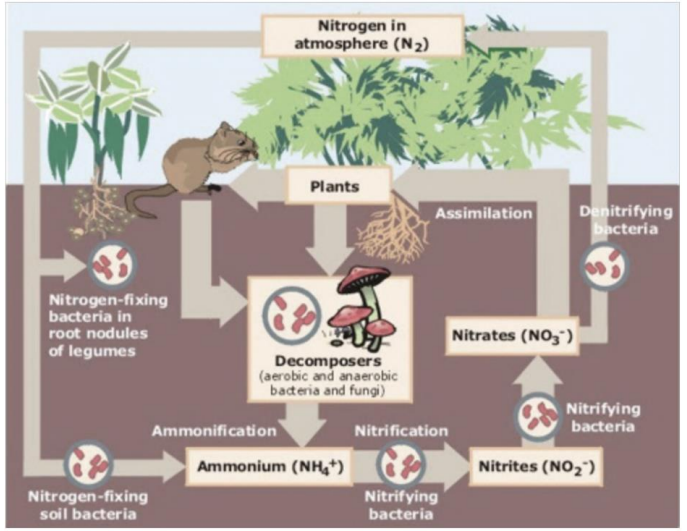

4. Open Research Directions

Agenda

1. Event process completion

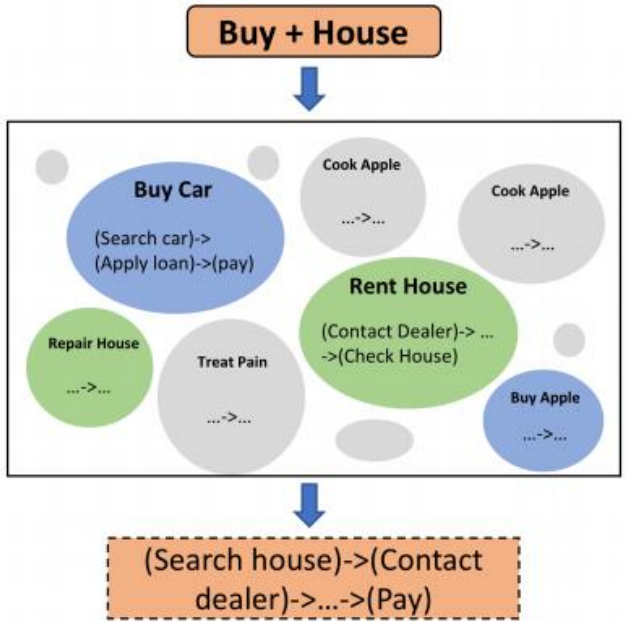

2. Event intention prediction

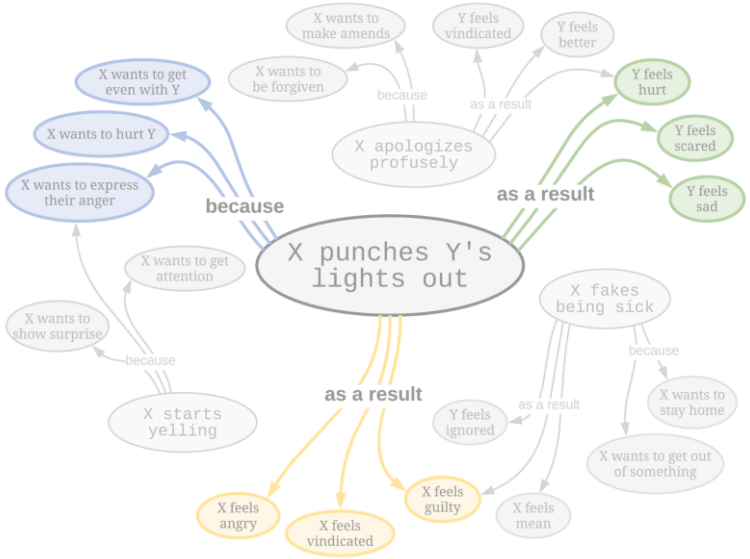

3. Event processes in downstream NLU tasks

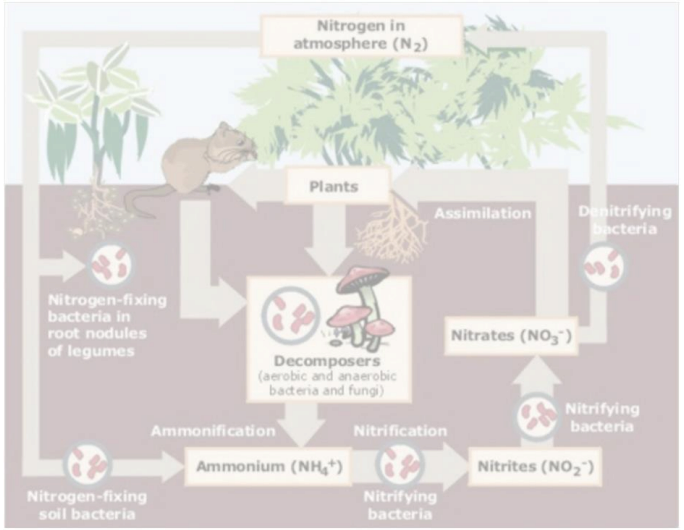

4. Open Research Directions

Event Process Completion

Two forms of process prediction

1. Predicting steps of the process

2. Inducing the entire process from scratch.

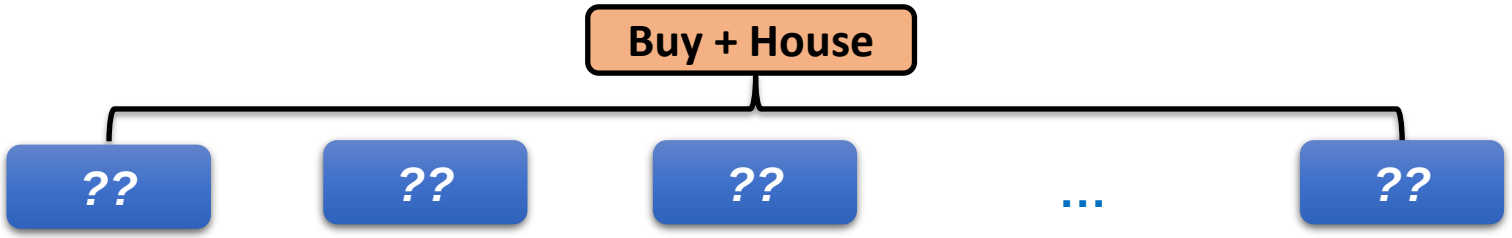

Event Process Completion

Chambers and Jurafsky. Unsupervised Learning of Narrative Event Chains. ACL-08

Unsupervised event process completion can be done using corpus statistics (Gigaword in this work)

- Capturing the co-occurrence of events using pointwise mutual information

$$pmi(e(w,d),e(v,g))$$

- The next most likely forthcoming event can be found by maximizing the accumulated PMI

$$\max_{j:0<j<m} \sum_{i=0}^{n} pmi(e_i, f_j)$$

(*n*: #events in the process; *m*: #events in the vocabulary.

Known events:
(pleaded subj), (admits subj), (convicted obj)

Likely Events:

sentenced obj	0.89	indicted obj	0.74
paroled obj	0.76	fined obj	0.73
fired obj	0.75	denied subj	0.73

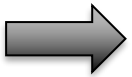

Improves narrative cloze tests (36% improvement on NYT Narrative Cloze).

Event Process Completion

Radinsky and Horvitz. Mining the Web to Predict Future Events. WSDM, 2013

Extension of the event chain model on multiple **dated** and **topically cohesive** documents.

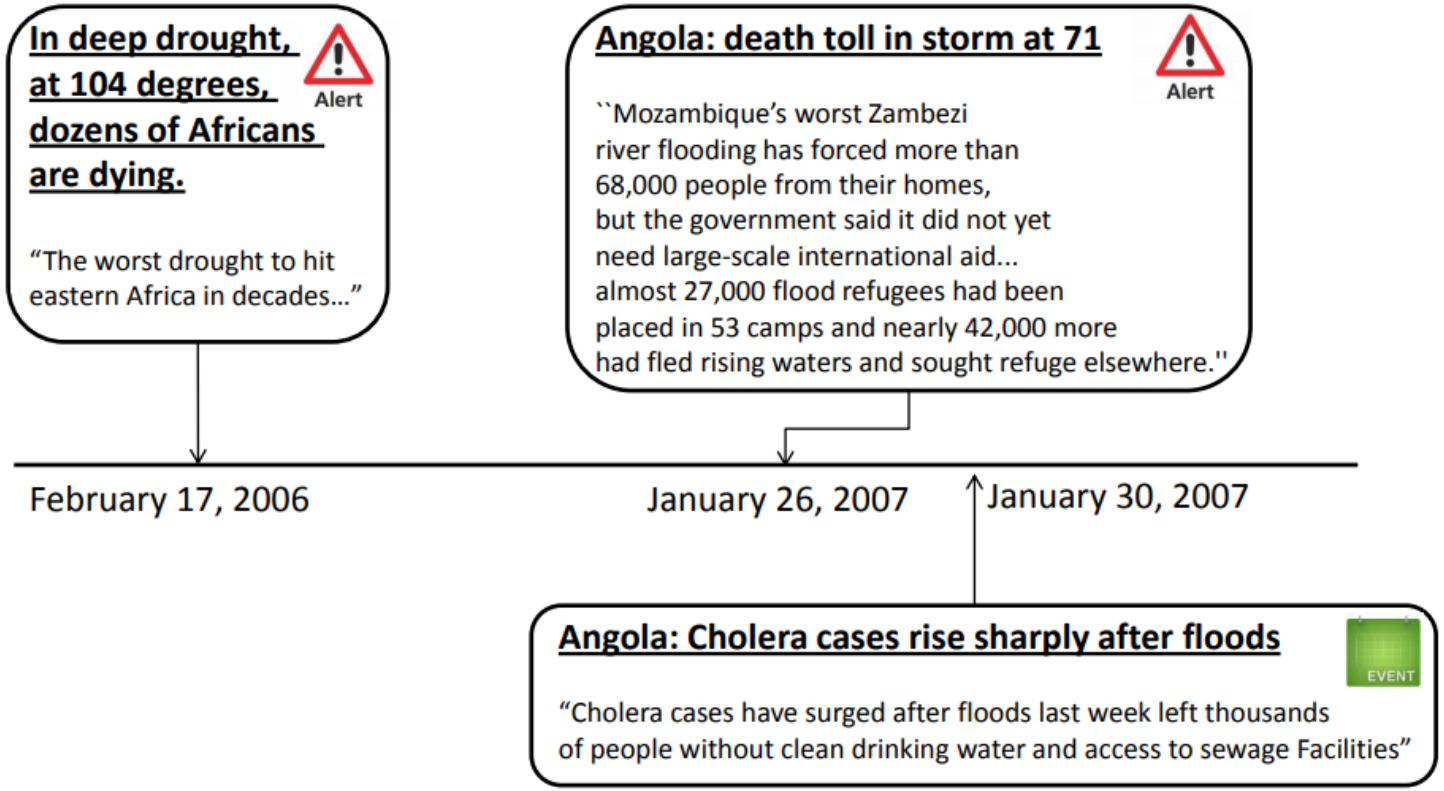

Maximum-entropy-based event co-occurence model.

The likelihood of cholera rising is predicted high after a drought followed by storms in Angola (*based on corpus statistics*).

Analogous Event Process Induction

Can we perform *de novo* event process induction?

Zhang, et al. Analogous Process Structure Induction for Sub-event Sequence Prediction. EMNLP, 2020

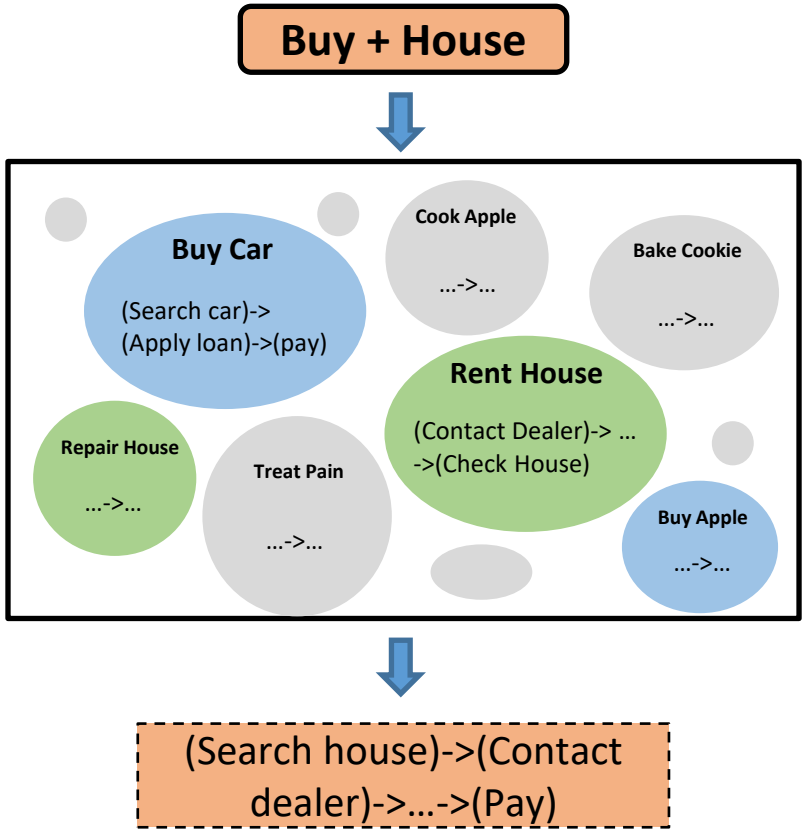

Analogous Event Process Induction

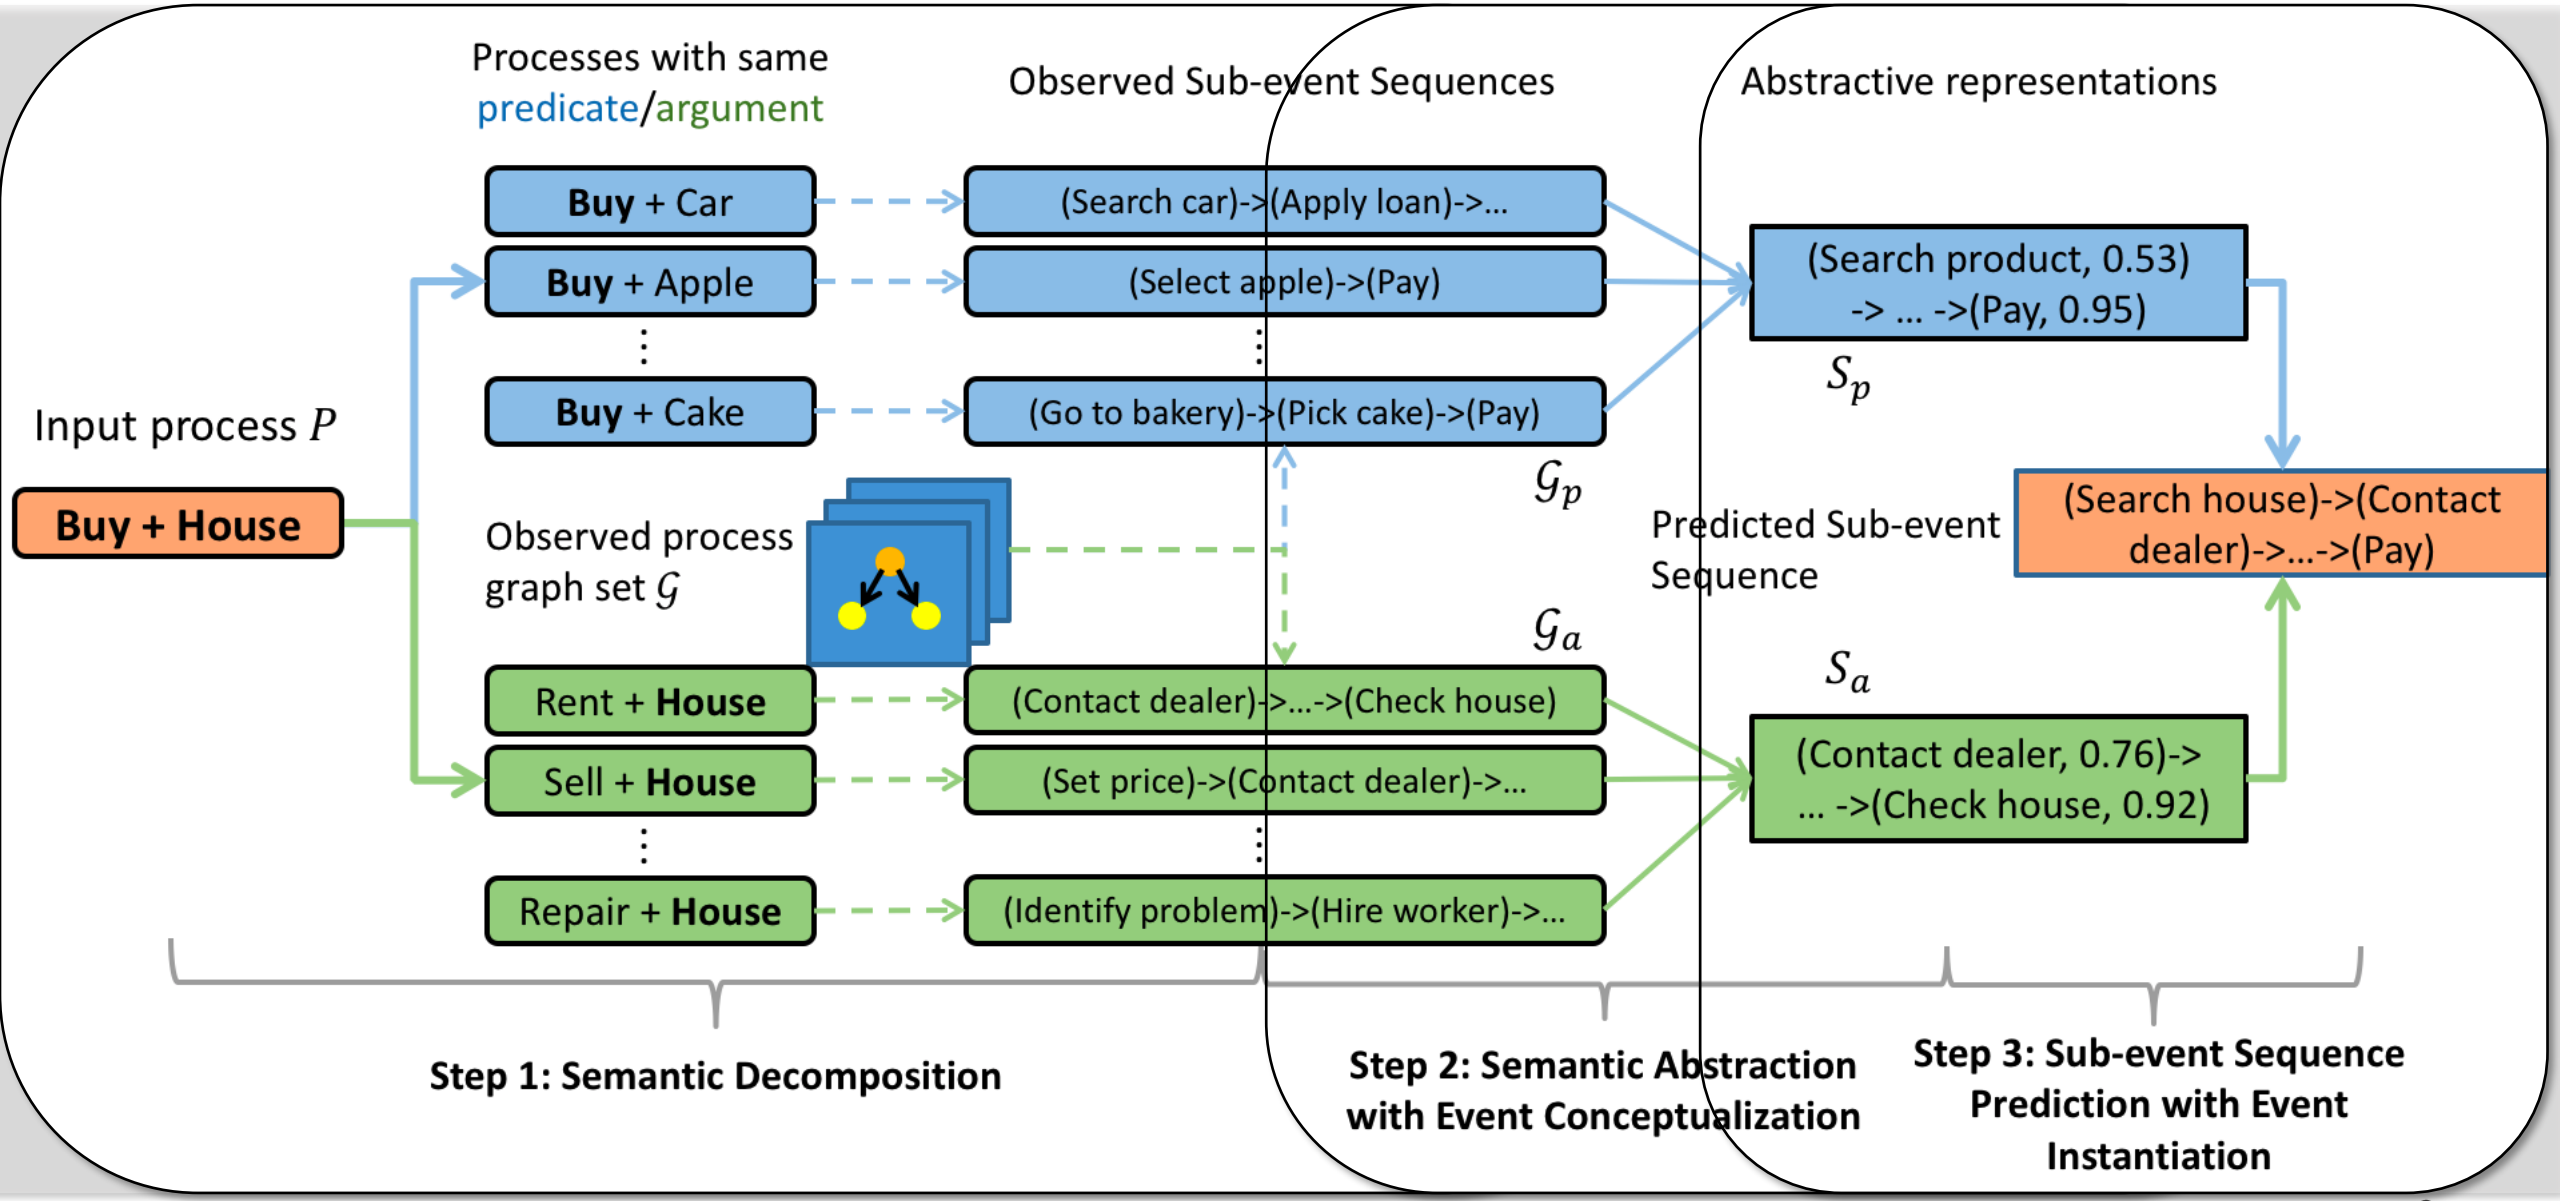

Evaluation Based on wikiHow Event Processes

Model	String Match E-ROUGE1	String Match E-ROUGE2	Hypernym Allowed E-ROUGE1	Hypernym Allowed E-ROUGE2
Random	2.9165	0.4664	23.5873	8.1089
Seq2seq (GloVe)	5.0323	1.4965	27.8710	13.0946
Seq2seq (RoBERTa)	4.5455	0.4831	28.0032	12.8502
Top one similar process (Jaccard)	8.8589	5.1000	28.6548	14.6231
Top one similar process (GloVe)	9.8797	5.1452	29.4203	13.6001
Top one similar process (RoBERTa)	9.2599	4.7390	30.6599	15.8417
Analogous Process Structure Induction (APSI)	**14.8013**	**6.6045**	**36.1648**	**19.2418**
Human	29.0189	15.2542	50.4647	29.4423

(a) Basic Setting (for each sub-event, we only predict and evaluate the verb)

Model	String Match E-ROUGE1	String Match E-ROUGE2	Hypernym Allowed E-ROUGE1	Hypernym Allowed E-ROUGE2
Random	0.0000	0.0000	0.5104	0.0903
Seq2seq (GloVe)	0.1935	0.0534	0.9677	0.1069
Seq2seq (RoBERTa)	0.4870	0.0000	1.7857	0.2899
Top one similar process (Jaccard)	0.6562	0.2257	2.4797	0.5867
Top one similar process (GloVe)	0.8750	0.2106	2.8801	0.7372
Top one similar process (RoBERTa)	0.9479	0.3009	3.2811	0.9929
Analogous Process Structure Induction (APSI)	**3.4988**	**0.4513**	**6.1611**	**1.1885**
Human	11.6351	5.5905	18.0034	8.2695

(b) Advanced Setting (for each sub-event, we predict and evaluate all words)

Process Name: **Treat Pain**

References: (‘learn cause’->‘identify symptom’->‘see doctor’)
(‘identify cause’->‘learn injury’->‘recognize symptom’->`recognize symptom’)

APSI Prediction: (‘Identify symptom’->‘see doctor’->‘recognize symptom’->‘take supplement’)

Quantitative results

Qualitative results

Resources are available at https://cogcomp.seas.upenn.edu/page/publication_view/910

Agenda

1. Event process completion

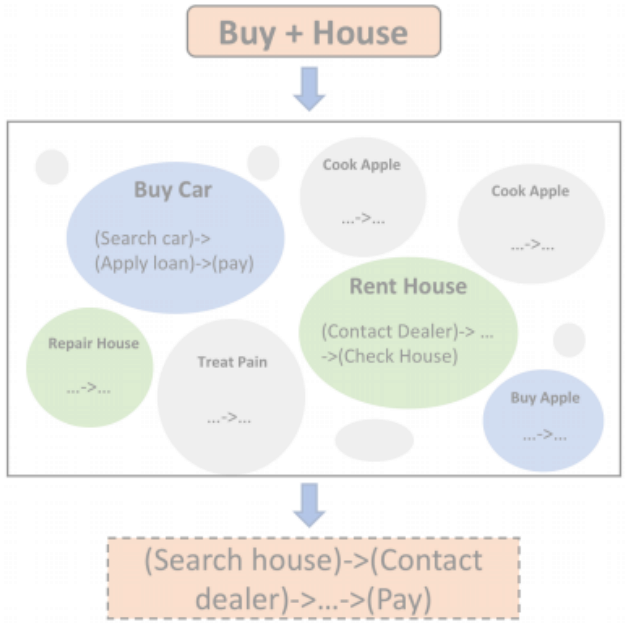

2. Event intention prediction

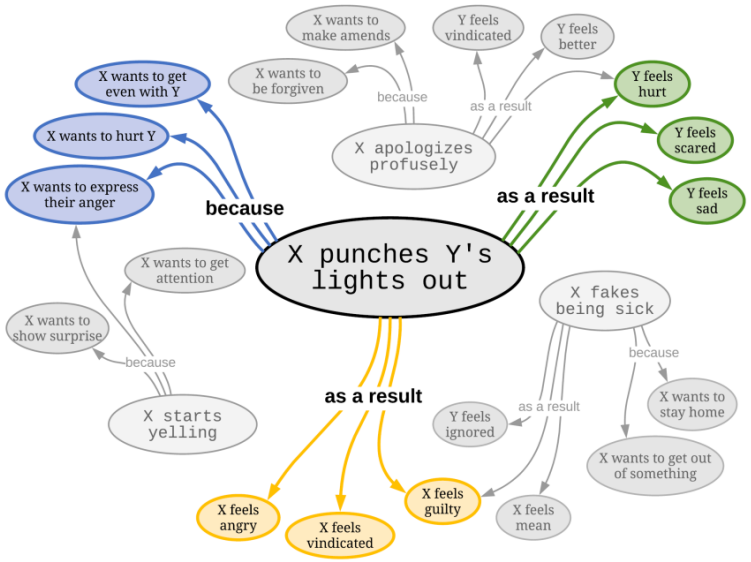

3. Event processes in downstream NLU tasks

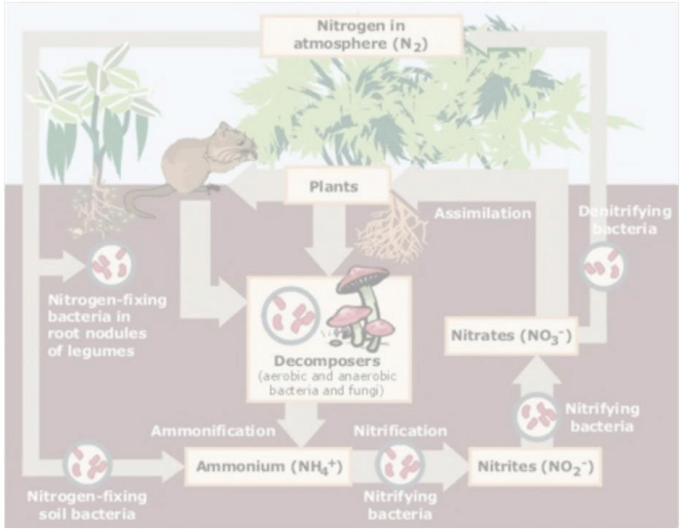

4. Open Research Directions

Intention Prediction for Events

People can easily anticipate the intents and possible reactions of participants in an event.

PersonX cooks thanksgiving dinner	X's intent	to impress their family
	X's reaction	tired, a sense of belonging
	Y's reaction	impressed
PersonX drags PersonX's feet	X's intent	to avoid doing things
	X's reaction	lazy, bored
	Y's reaction	frustrated, impatient

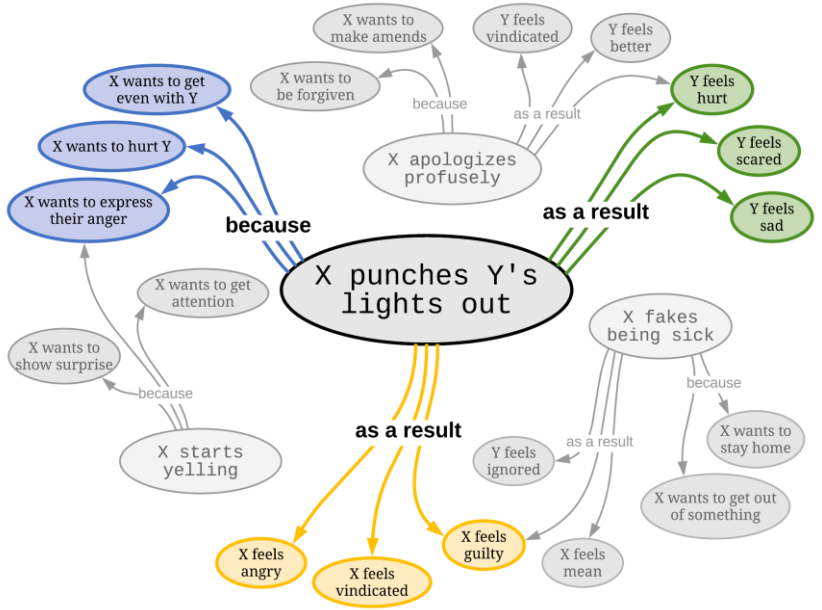

A commonsense-aware system should also perform such prediction.

Event2Mind – A learning system that understands stereotypical intents and reactions to events (Rashkin et al. ACL-18)

Event2Mind

Is developed based on large crowdsourced corpora:

- 25,000 events
- Free-form descriptions of their intents and reactions

Performs Seq2NGram generation:

PersonX cooks steak

PersonX's intent: ["steak", "to kill their hunger", "to make dinner for the family", "to eat steak"]

PersonX's reaction: ["excited", "accomplished", "proud", "full"]

Other people's reaction: ["none", "happy", "person x cooked well."]

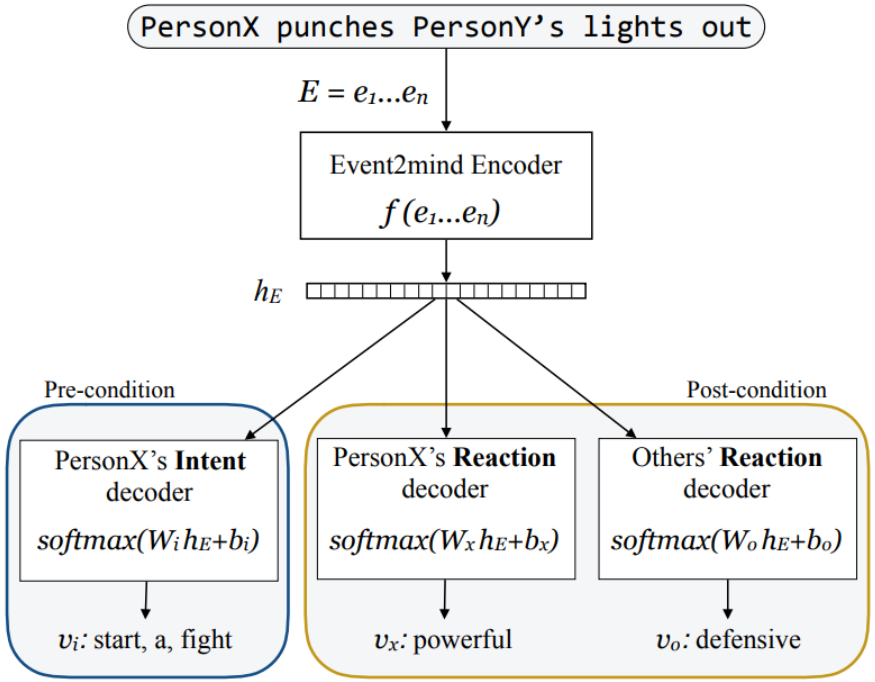

More follow-ups of Event2Mind

- ATOMIC: An Atlas of Machine Commonsense for If-Then Reasoning (Sap+ AAAI 2019)
- COMET: Commonsense Transformers for Automatic Knowledge Graph Construction (Bosselut+, ACL-19)

Intention Prediction for Event Processes

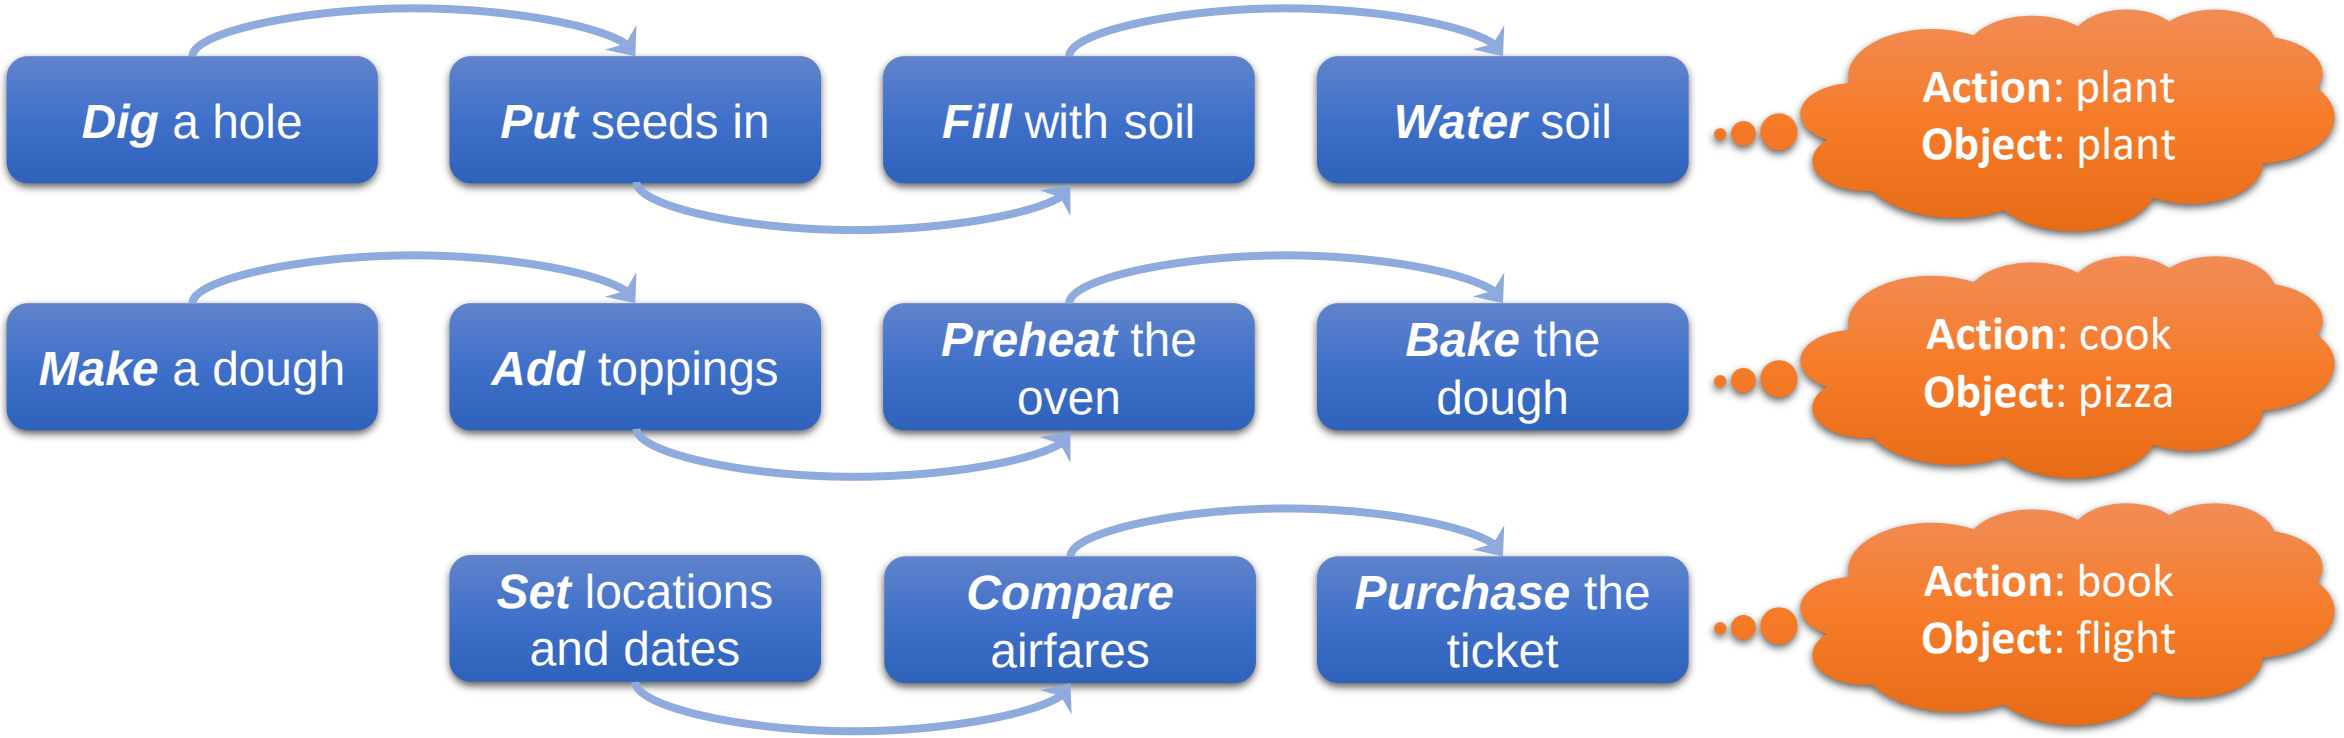

Event processes are directed by the **central goal**, or the **intention** of its performer [Zacks+, Nature Neuroscience 2001].

- Inherent to human's common sense.
- Missing from current computational methods.
- Important to machine commonsense reasoning, summarization, schema induction, etc.

A New Task: Multi-axis Event Process Typing

Chen et al. “What are you trying to do?” Semantic Typing of Event Processes. CoNLL-2020 **(Best Paper Nomination)**

A new (cognitively motivated) **semantic typing task** for understanding event processes in natural language. Two **type axes**:

- What ***action*** the event process seeks to take? (**action type**)
- What type of **object**(s) it should affect? (**object type**)

This research also contributes with

- A **large dataset** of typed event processes (>60k processes)
- A **hybrid learning framework** for event process typing based on **indirect supervision**

A Large Event Process Typing Dataset

A large dataset of typed event processes from wikiHow

- 60,277 event processes with free-form labels of action and object types

A challenging typing system

- **Diversity**: 1,336 action types and 10,441 object types (in free froms)
- **Few-shot cases**: 85.9% labels appear less than 10 times, (~half 1-shot).
- **External labels**: in 91.2% (84.2%) processes, the action (object) type label does not appear in the process body.

A non-trivial learning problem with **ultra fine-grained** and **extremely few-shot** labels.

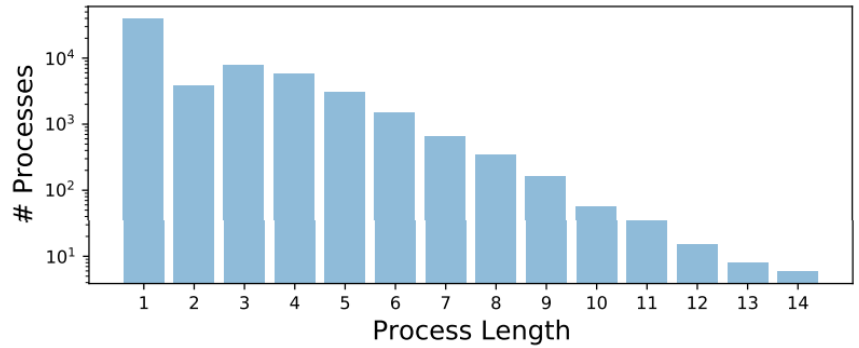

Figure 2: Distribution of process lengths.

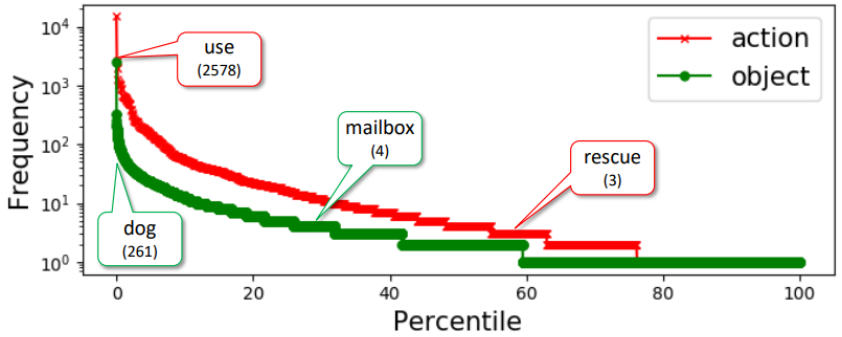

Figure 3: Distribution of actions and objects. Number of frequencies are shown in the brackets.

Indirect Supervision from Gloss Knowledge

An event process

Directly inference
(Difficult)

Make
Cocktail
Labels

Indirect inference
(Much Easier)

Make	create or manufacture a man-made product
Cocktail	a short, mixed drink

Label glosses (from WordNet)

Why using label glosses?

- Semantically richer than labels themselves
- Capturing the association of a process-gloss pair (two sequences) is much easier
- Jump-starting few-shot label representations (and benefiting with fairer prediction)

Indirect Supervision from Gloss Knowledge

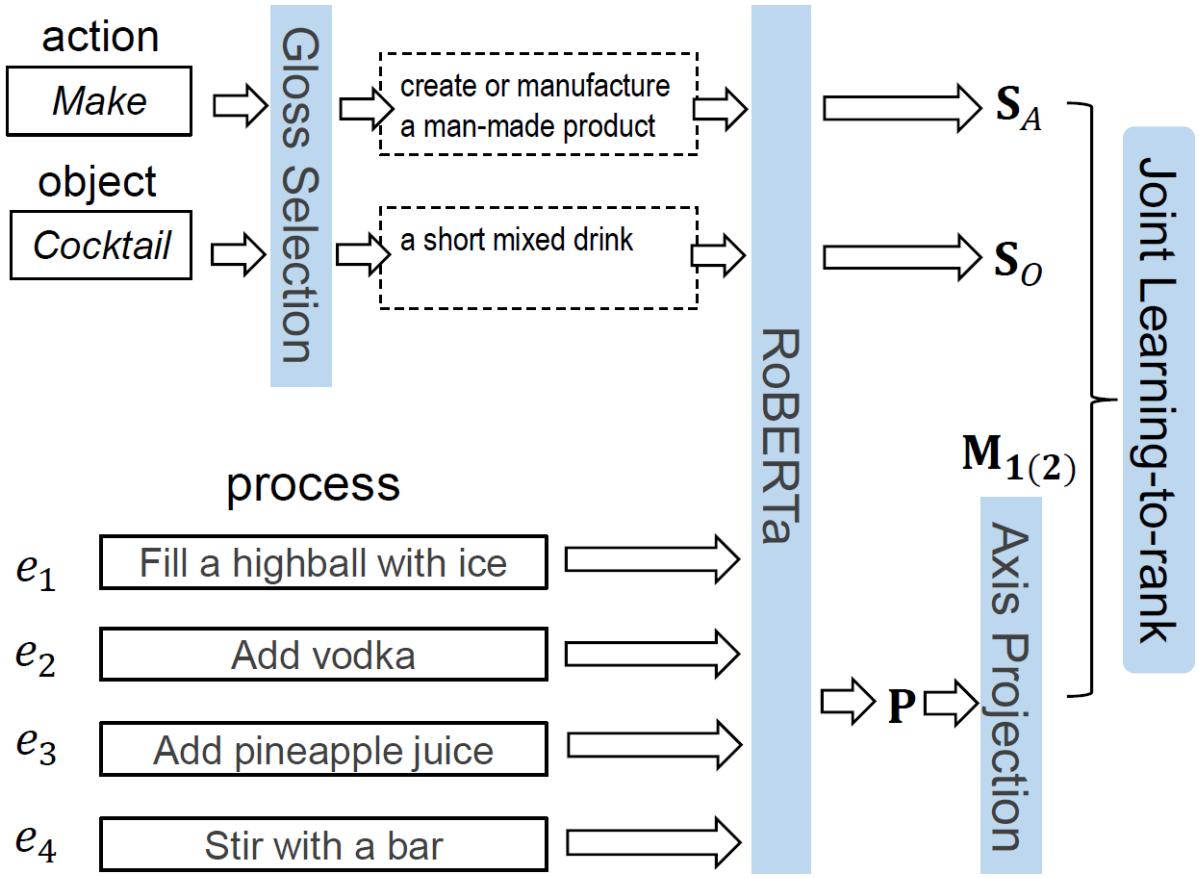

How to represent the process?

- RoBERTa encodes concatenated event contents (VERB and ARG1).

How to represent a label?

- The same RoBERTa encodes the label gloss

Which gloss for a polysemous label?

- WSD [Hadiwinoto+, EMNLP-19]
- MFS (Most frequent sense)

Learning objective?

- Joint **learning-to-rank** for both type axes (different projection)

Inference?

- Ranking all glosses for all labels in the vocab

Results

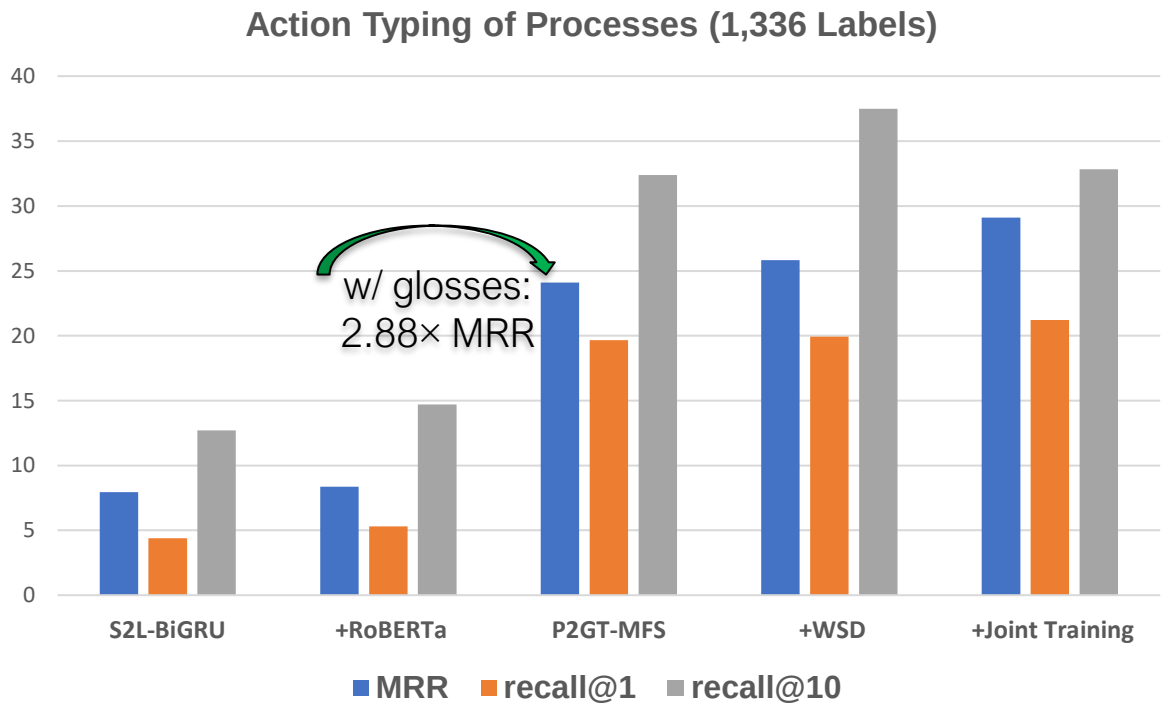

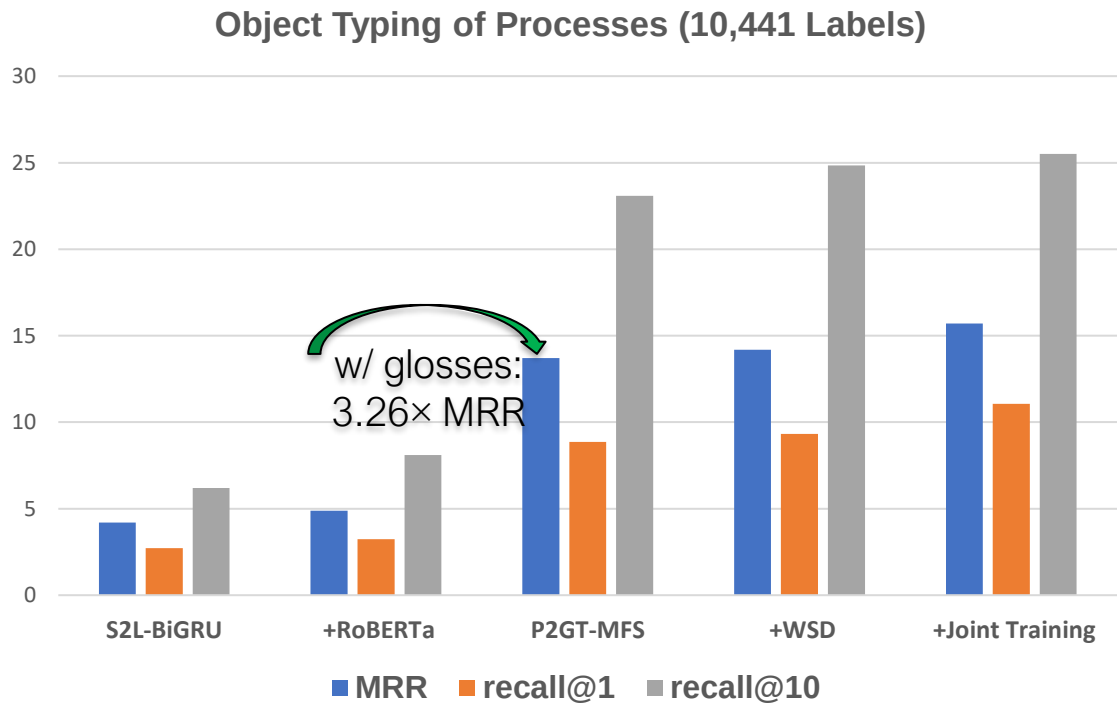

- Gloss knowledge brings along the most improvement (2.88~3.26 folds of MRR)
- Joint training indicates the effectiveness of leveraging complementary supervision signals
- Sense selection (WSD) leads to lesser improvement (predominant senses are representative enough)

Case Study

Event processes	Predictions
Make explosive materials ⇒ Obtain a container ⇒ Obtain shrapnel ⇒ Install a trigger	*A*: **detonate**, **assemble**, blacken *O*: **grenade**, **blaster**, **mine**
Go to DMV ⇒ Take photos ⇒ Take vision test ⇒ Take permit test ⇒ Take road test	*A*: **obtain**, *verify*, explore *O*: **license**, check, visa
Ignore order ⇒ Enter area ⇒ Enforce blockade ⇒ Force to retreat from area	*A*: **conquer**, *disarm*, **invade** *O*: *barrier*, **soldier**, **fortress**
Capture two opposition posts ⇒ Kill many fighters ⇒ Destroy three armed trucks ⇒ Confiscate artillery guns	*A*: *kill*, **demolish**, **fight** *O*: *melee*, **conflict**, **stronghold**
Cooperate with the counsel investigation ⇒ Open his remarks ⇒ Apologize many times ⇒ Try to restore public trust	*A*: *respond*, disagree, **accept** *O*: *apology*, *disagreement*, **slander**
Travel in a presidential motorcade ⇒ Be shot once in the back ⇒ Be taken to hospital ⇒ Be pronounced dead	*A*: *survive*, **die**, tackle *O*: **assassin**, crash, *roadkill*
Give advance notice ⇒ Give notice ⇒ Issue dividends	*A*: **honor**, **pay**, **reward** *O*: *finance*, **equity**, **subsidy**
Target quotes ⇒ Target shares quotes ⇒ Ask to clarify offer ⇒ Challenge to merge agreement ⇒ Challenge to merge businesses	*A*: **compare**, **maximize**, **negotiate** *O*: *prospectus*, **quote**, **settlement**
Clean windows ⇒ Buy plants ⇒ Hang pictures ⇒ Paint walls ⇒ Carpet floors	*A*: **redecorate**, **decorate**, *refurbish* *O*: **room**, **bedroom**, *makeover*

Table 3: Case study for typing event processes in the news domain. The predictions are given by Joint P2GT-WSD trained on our full dataset. Each case is given top 3 predictions on both axes, whereof reasonably correct ones are boldfaced, and relevant ones are italic. Few-shot labels appearing up to 10 times in our dataset are in blue.

System Demonstration

A web demonstration of our prototype system is running at https://cogcomp.seas.upenn.edu/page/demo_view/step

Examples

Decoration

Event process (choose an example or write the subevents of a process separated by '@' to get its intention)

clean windows @ buy plants @ paint walls @ hang pictures @ carpet floors @ reorganize furniture

Get intention >

redecorate room

Cosine similarity	Action	Object	Cosine similarity
0.678	redecorate	room	0.623
0.650	stage	atmosphere	0.599
0.500	brighten	mosaic	0.589
0.427	preoccupy	suite	0.574
0.418	furnish	interior	0.573

Agenda

1. Event process completion

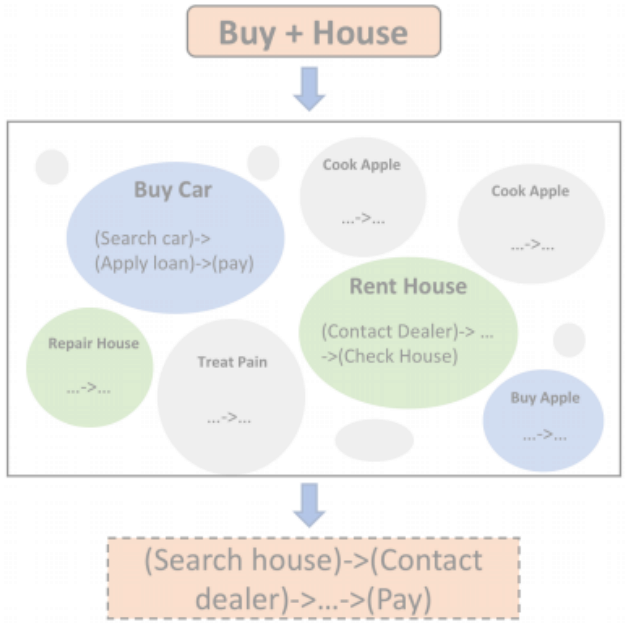

2. Event intention prediction

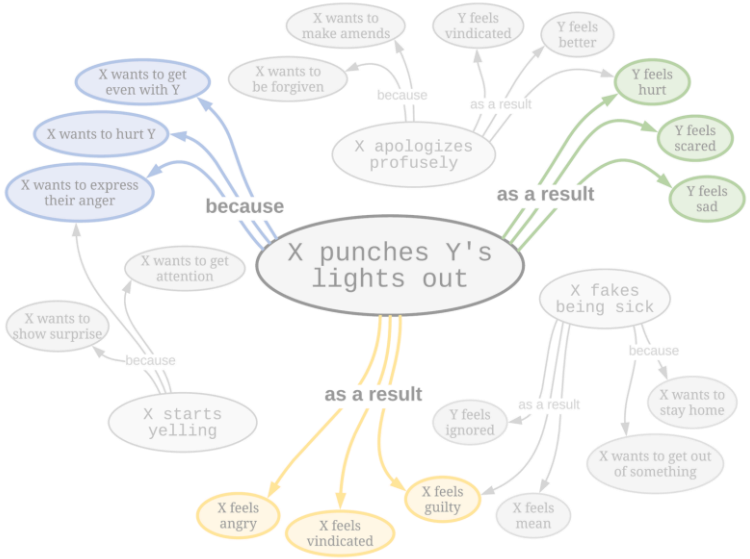

3. Event processes in downstream NLU tasks

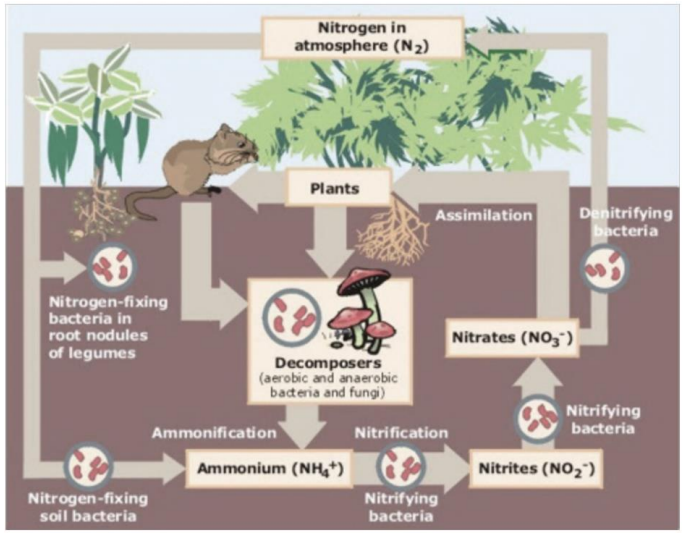

4. Open Research Directions

Narrative Prediction

The ROC Story Narrative Cloze Test [Mostafazadeh+, NAACL 2016]**:**

> One day Wesley's auntie came over to visit. He was happy to see her, because he liked to play with her. When she started to give his little sister attention, he got **jealous**. He got **angry** at his auntie and **bit** his sister's hand when she wasn't looking.

Then what might happen?

O1: He was **scolded**.

O2: She **gave him a cookie** for being so nice.

Chaturvedi, et al (EMNLP, 2017) train a language model that captures three types of sequential features:

1. **Event sequences** in 20 years of NYT data

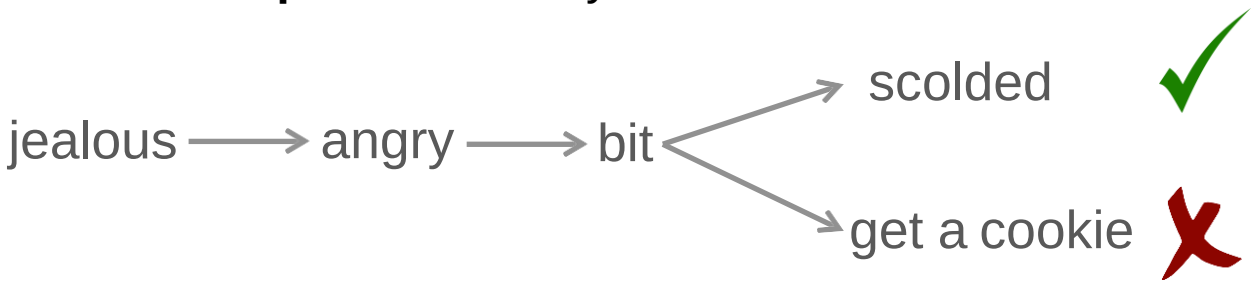

2. Sentiment trajectories
3. Topical consistency

Features	Accuracy
All	74.4%
Event-sequence	71.6%
Sentiment	64.5%
Topic	55.2%

Event sequences are most important.

Machine Reading Comprehension

QA based on articles in biology

Water is split, providing a source of electrons and protons (hydrogen ions, H^+) and giving off O_2 as a by-product. ***Light absorbed*** by chlorophyll drives a ***transfer of the electrons and hydrogen ions*** from water to an acceptor called $NADP^+$.

What can the splitting of water lead to?

A: Light absorption **B:** Transfer of ions

1. Extracting events and event-event relations from articles

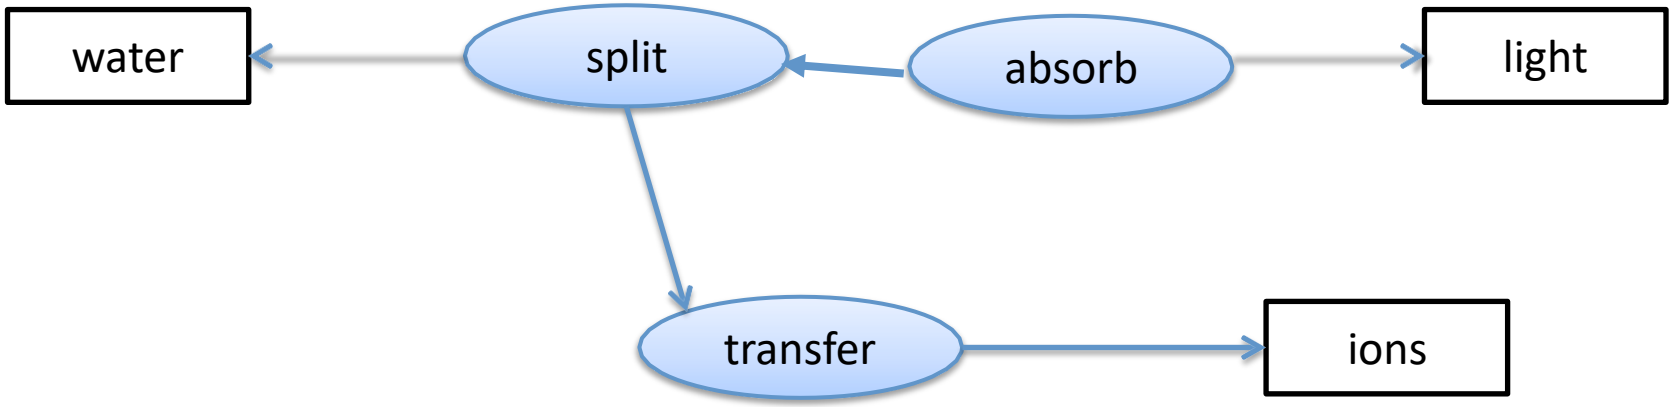

2. Matching questions and candidate answers with extracted event processes

Berant, et al. Modeling Biological Processes for Reading Comprehension. EMNLP, 2014 (Best Paper Award)

Video Segmentation

Events in a process as anchors of video segments.

wikiHow process: make pancakes: {add egg, add flour, …, pour batter, remove pancake}

Video segments:

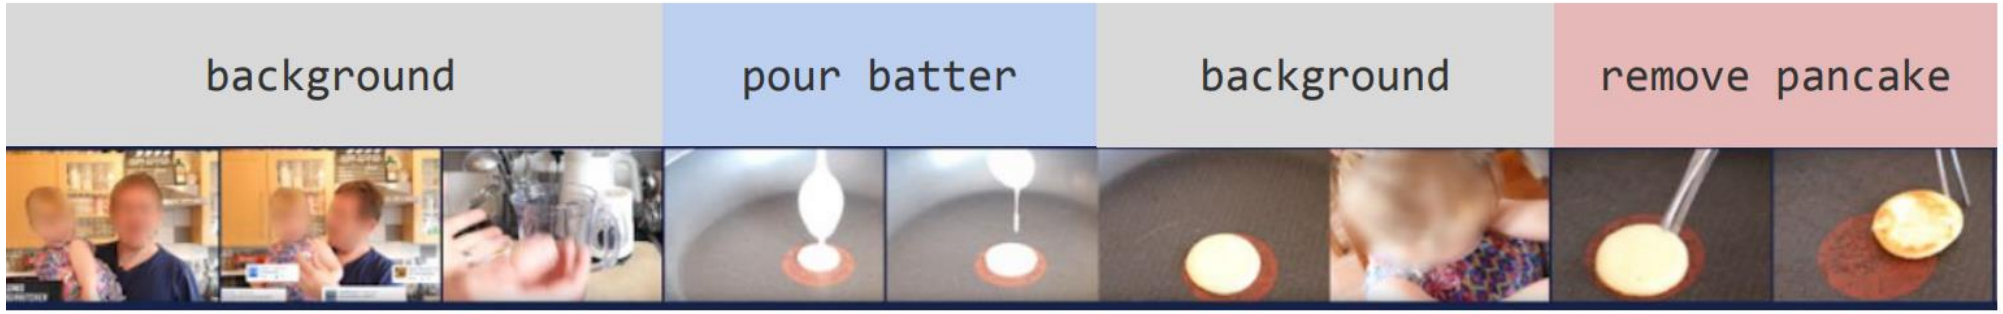

Video:

Video narration: *hey folks here welcome to my kitchen … pour a nice-sized amount … change the angle to show … and take it out*

Alignment learning between video narration and wikiHow event processes help action segmentation in videos.

Zhukov et al. Cross-task weakly supervised learning from instructional videos. CVPR 2019
Fried et al. Learning to Segment Actions from Observation and Narration. ACL 2020

Future Event Prediction in Videos

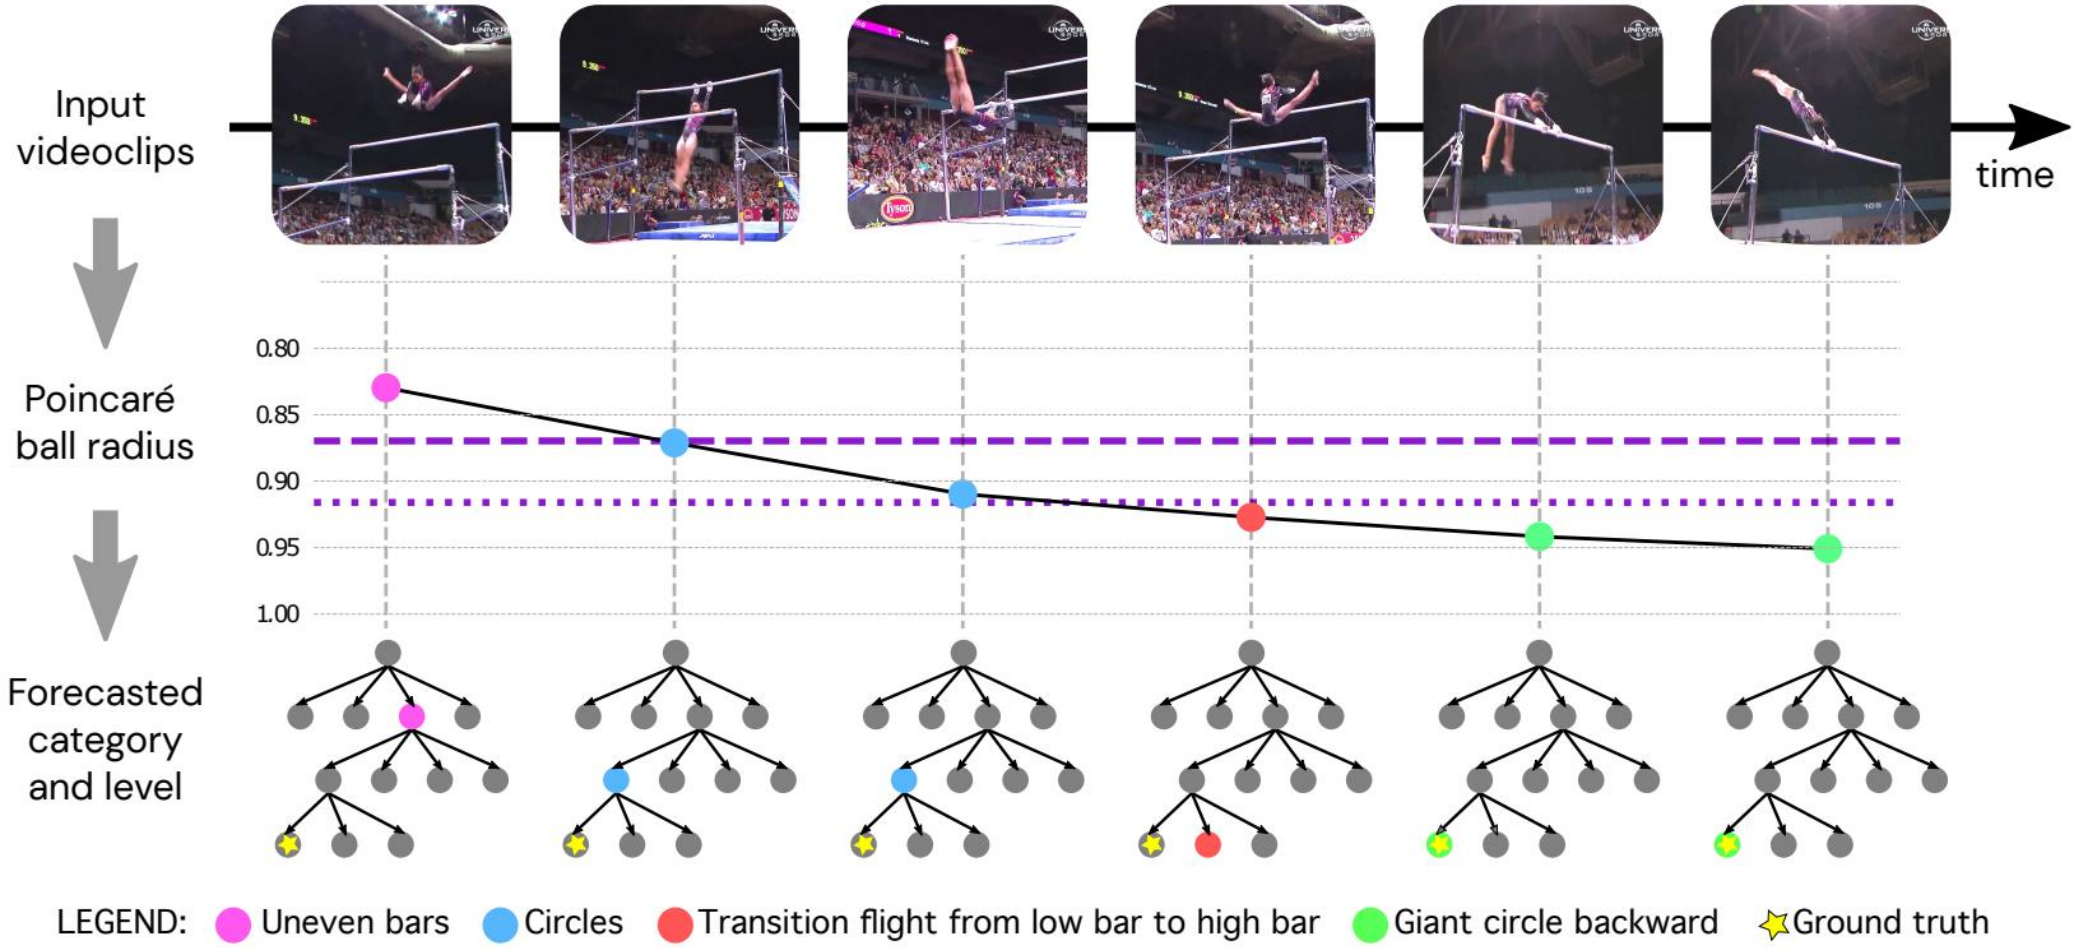

Hyperbolic embeddings model hierarchies of possible event evolution processes in videos.

Surís et al. Learning the Predictability of the Future. CVPR, 2021

Agenda

1. Event process completion

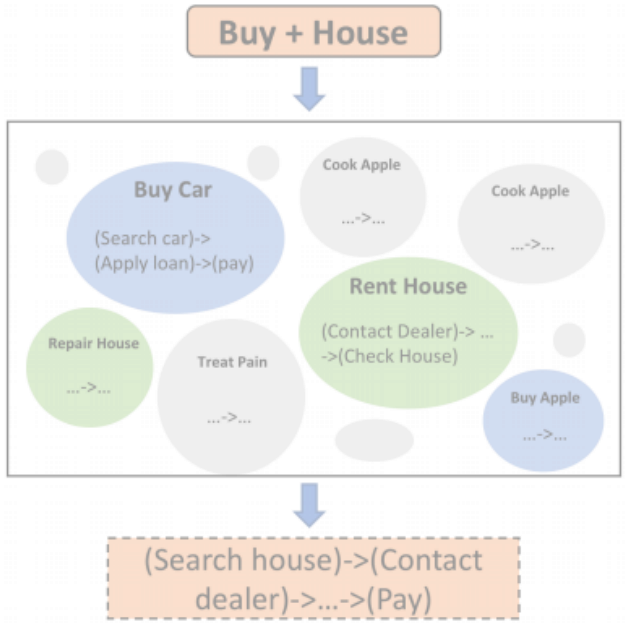

2. Event intention prediction

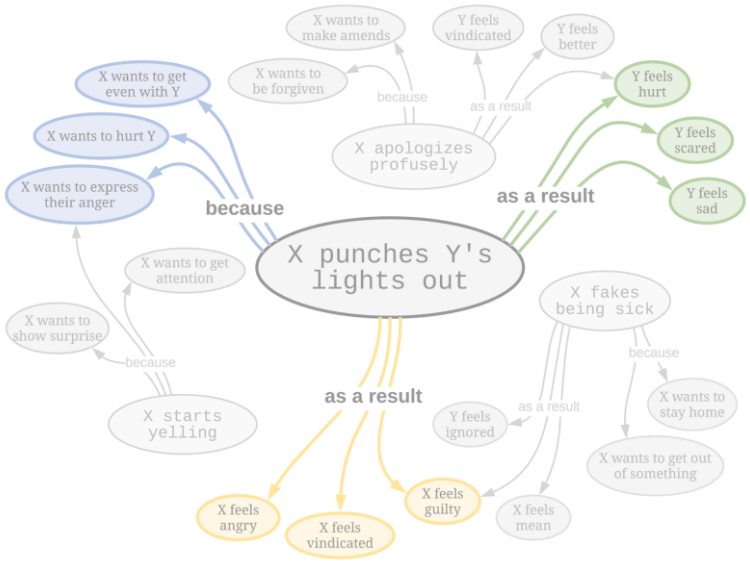

3. Event processes in downstream NLU tasks

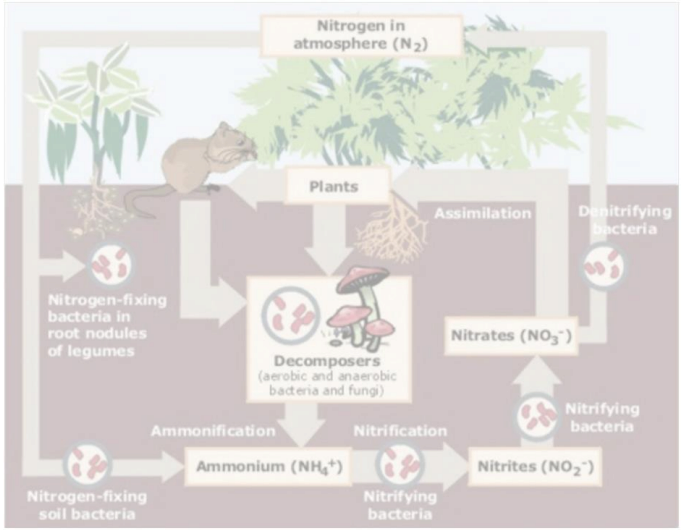

4. Open Research Directions

Salience/Essentiality Detection in Event Processes

Events in a process are not equally important

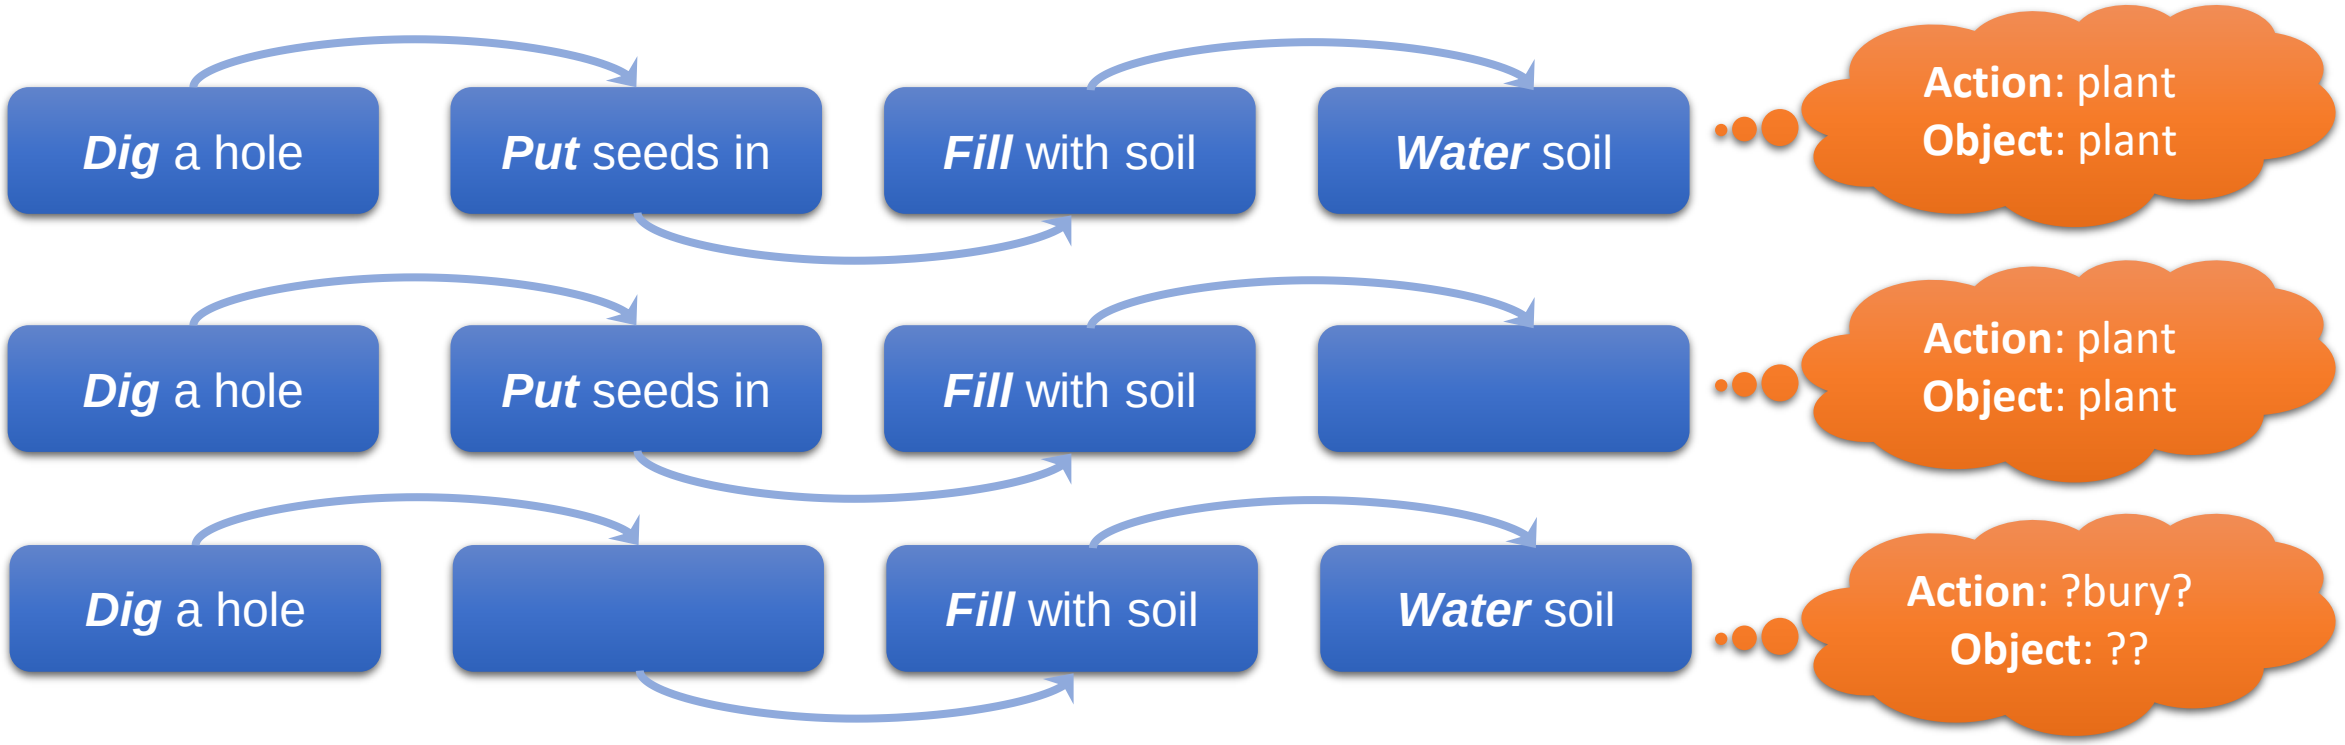

Defending your dissertation is essential; **Doing a TAShip** is less important; **Doing an internship** is optional…

How to automatically identify salient events in a process?
Would those help downstream tasks such as abstractive summarization?

(Temporal) Commonsense Understanding

Do language models understand:

Time duration

- Earning a PhD takes **several years**; not **several months**; not **lifelong time**.
- Having a banquet dinner **takes around an hour**; not **several minutes**; not **a day**.

Typical time

- People eat break fast in the morning.
- Tornados may strike Florida typically in the middle of a year.

Typical frequency

- Cars change oil every **year/half year.**
- People pay utility bills every **months/two months**.

Ben Zhou and Daniel Khashabi and Qiang Ning and Dan Roth. *"Going on a vacation" takes longer than "Going for a walk": A Study of Temporal Commonsense Understanding",* **EMNLP** 2019
Ben Zhou and Qiang Ning and Daniel Khashabi and Dan Roth. *Temporal Common Sense Acquisition with Minimal Supervision*, **ACL** 2020.

Reasoning About Event Ordering

Identifying the order of member events in a process?

Heavy snow is causing disruption to transport across the UK, with heavy rainfall bringing flooding to the south-west of England. Rescuers searching for a woman trapped in a landslide at her home in Looe, Cornwall, said they had found a body.

Q1: What events have already finished?
A: searching trapped landslide said found
Q2: What events have begun but has not finished?
A: snow causing disruption rainfall bringing flooding
Q3: What will happen in the future?
A: No answers.

warm-up

Q4: What happened before a woman was trapped?
A: landslide
Q5: What had started before a woman was trapped?
A: snow rainfall landslide
Q6: What happened while a woman was trapped?
A: searching
Q7: What happened after a woman was trapped?
A: searching said found

User-provided

Ning, et al. TORQUE: A Reading Comprehension Dataset of Temporal Ordering Questions. EMNLP, 2020

- 3.2k news snippets with 21k human-generated questions querying temporal relationships

Lyu, et al. Reasoning about Goals, Steps, and Temporal Ordering with WikiHow. EMNLP, 2020

- A wikiHow-based testbed about event ordering (and more)

Constrained story generation based on events?

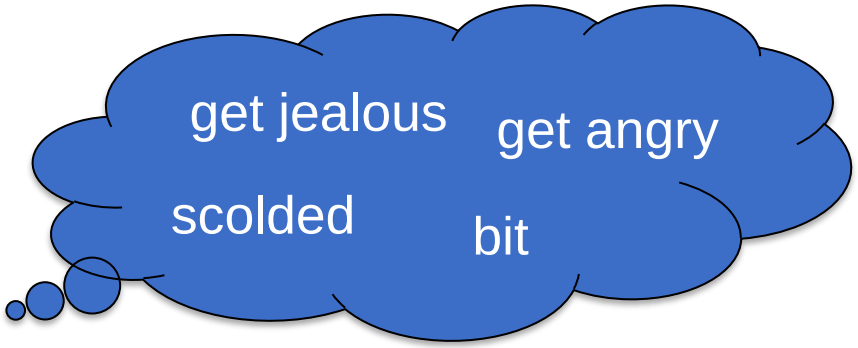

jealous → angry → bit → scolded

He got **jealous**. He got **angry** at his auntie and **bit** his sister's hand when she wasn't looking. Then he was **scolded**.

“Special Cases” of Event Extraction

Unsupervised Event Extraction

Yesterday, **Jeff *shaved* my hair** at home.

Semantic Role Labeling

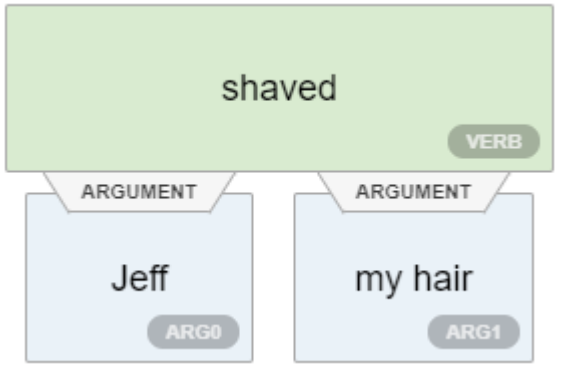

What about:

Nominal events? Tournament, War, Match, …

Can we detect them based on word senses?

Imaginary events? Jeff planned to **shave his hair** yesterday, but he was too busy to do that.

Verbs that are not event triggers? Jeff’s haircut **looks** good … Fresh air **smells** good.

More Tasks

Chatbots

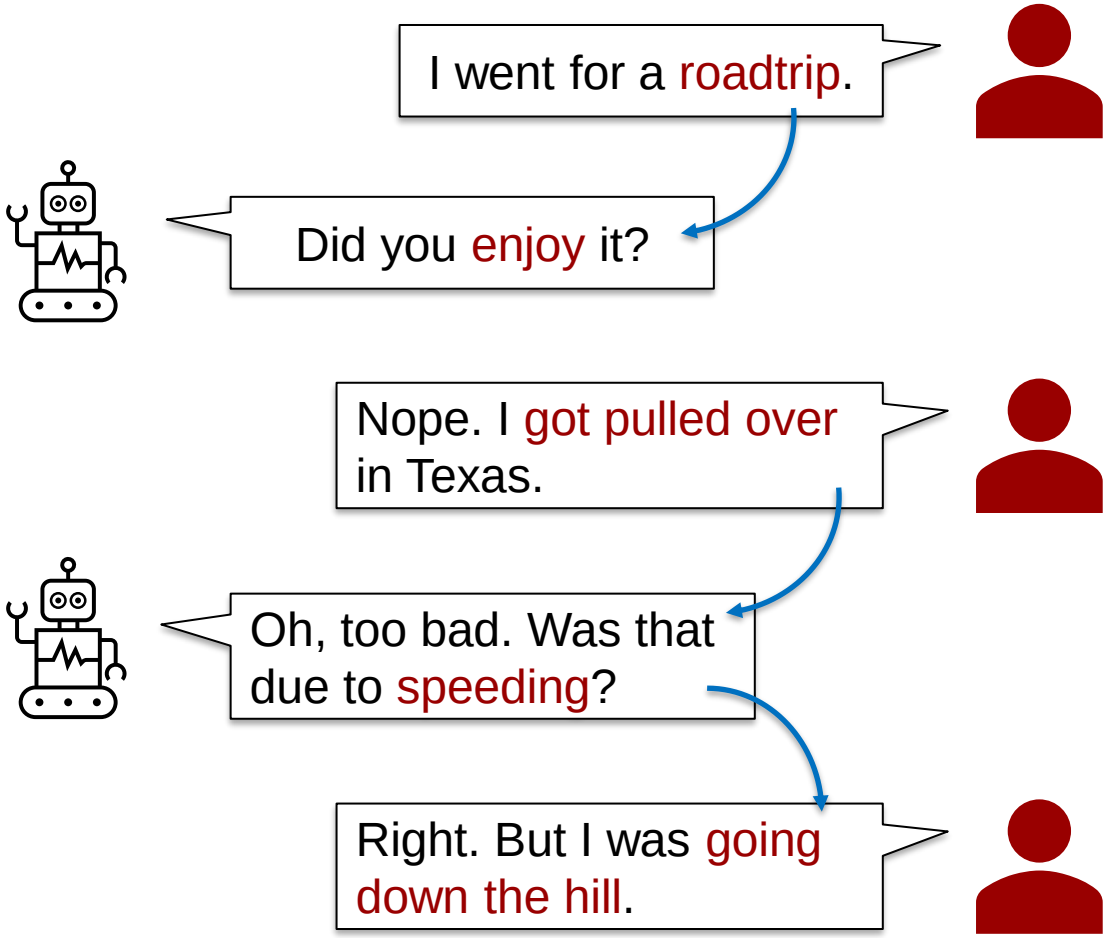

Can event processes improve the consistency of utterance generation/retrieval?

Understanding clinical event processes

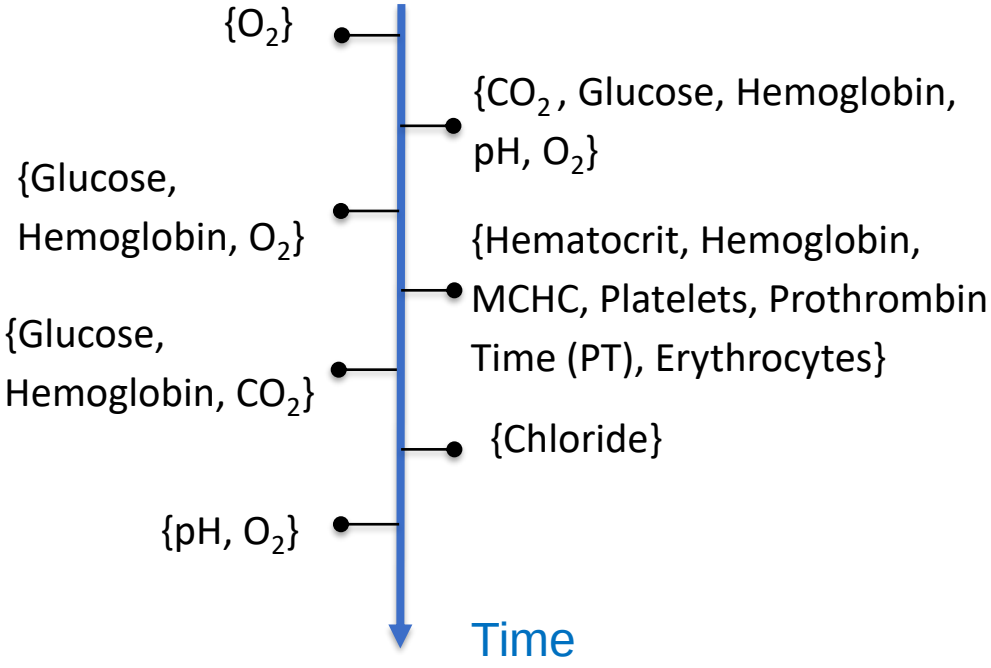

Diagnostic prediction (Zhang et al. AIME-20), phenotype prediction, …

- Transfer learning can be important (naturally lack of data)
- Structured prediction can be important (dependency of phenotypes, disease labels)

COLUMBIA Penn USC Viterbi School of Engineering UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN amazon

Event-Centric Natural Language Processing

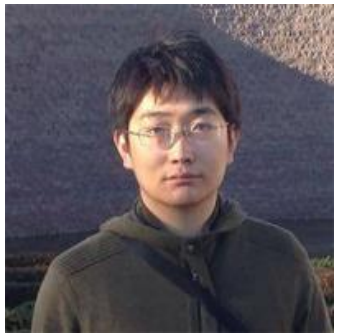
Muhao Chen

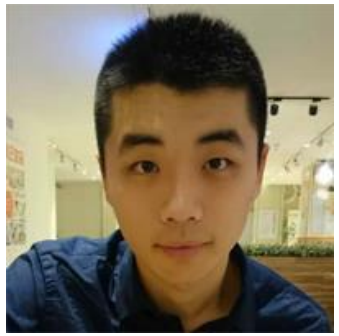
Hongming Zhang

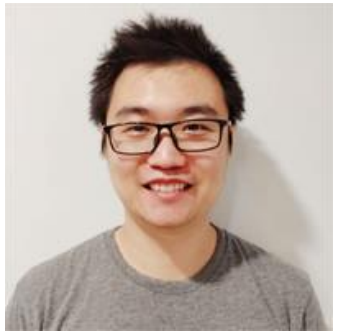
Qiang Ning


Manling Li


Heng Ji


Kathleen McKeown


Dan Roth

Agenda

- **Event extraction** (Manling & Heng @UIUC)
- **Event relation extraction** (Qiang @Amazon)
- **Event process understanding** (Muhao @USC)
- **Eventuality knowledge acquisition** (Hongming @UPenn)
- **Event Summarization** (Kathleen @Columbia)
- **The future of event-centric NLP** (Dan @UPenn)

Association for Computational Linguistics

ACL-IJCNLP 2021

The **Event-Centric Natural Language Processing** Tutorial

- At **ACL, August 2021**

References

- Zack, et al. Human brain activity time-locked to perceptual event boundaries. Nature neuroscience, 4(6):651–655. 2001
- Chambers and Jurafsky. Unsupervised learning of narrative event chains. ACL, 2008
- Radinsky and Horvitz. Mining the Web to Predict Future Events. WSDM, 2013
- Berant, et al. Modeling Biological Processes for Reading Comprehension. EMNLP, 2014
- **Chaturvedi, et al. Story comprehension for predicting what happens next. EMNLP, 2017**
- Rashkin, et al. Event2Mind: Commonsense Inference on Events, Intents, and Reactions. ACL, 2018
- Liu, et al. Automatic event salience identification. EMNLP, 2018
- Zhukov et al. Cross-task weakly supervised learning from instructional videos. CVPR, 2019
- **Zhang, et al. Analogous Process Structure Induction for Sub-event Sequence Prediction. EMNLP, 2020**
- **Chen, et al. "What are you trying to do?" Semantic typing of event processes. CoNLL, 2020**
- **Ning et al. TORQUE: A Reading Comprehension Dataset of Temporal Ordering Questions. EMNLP, 2020**
- Lyu, et al. Reasoning about Goals, Steps, and Temporal Ordering with WikiHow. EMNLP, 2020
- **Jindai, et al. Is Killed More Significant than Fled? A Contextual Model for Salient Event Detection. COLING, 2020**
- Fried, et al. Learning to Segment Actions from Observation and Narration. ACL, 2020
- **Zhang, et al. Diagnostic Prediction with Sequence-of-sets Representation Learning for Clinical Events. AIME, 2020**
- Surís, et al. Learning the Predictability of the Future. CVPR, 2021

USC Viterbi School of Engineering

Penn University of Pennsylvania

ILLINOIS

amazon alexa

Thank You

Muhao Chen. Homepage: https://muhaochen.github.io/ Email: muhaoche@usc.edu

06/2021