清华大学人工智能研究院
知识智能研究中心
Knowledge Intelligence Research Center

常识知识感知的语言生成初探
Language Generation with Commonsense Knowledge

黄民烈

清华大学人工智能研究院
清华大学计算机系

Knowledge Everywhere

- **Knowledge type**
  - World facts
  - Commonsense knowledge
- Encoding **symbolic knowledge** becomes a hot topic
- **Application**
  - Language inference, semantic reasoning
  - MRC, QA & dialogue
  - Language generation (story, dialogue, etc.)

Commonsense Knowledge

- **Commonsense knowledge** consists of facts about the everyday world, that all humans are expected to know. (Wikipedia)
  - Lemons are sour
  - Tree has leaves
  - Dog has four legs
- Commonsense Reasoning ~ **Winograd Schema Challenge**:
  - The trophy would not fit in the brown suitcase because it was too **big**. What was too **big**?
  - The trophy would not fit in the brown suitcase because it was too ***small***. What was too ***small***?

Commonsense Extraction

- What is commonsense knowledge?
- What is the **boundary**?
- Commonsense extraction
  - From embeddings [1]
  - Commonsense knowledge base completion [2]
  - From raw data (text, image) [3]

① Yang et al. 2018. Extracting Commonsense Properties from Embeddings with Limited Human Guidance
② Li et al. 2018. Commonsense Knowledge Base Completion
③ Xu et al. 2018. Automatic Extraction of Commonsense LocatedNear Knowledge

CS Knowledge in Reading Comprehension

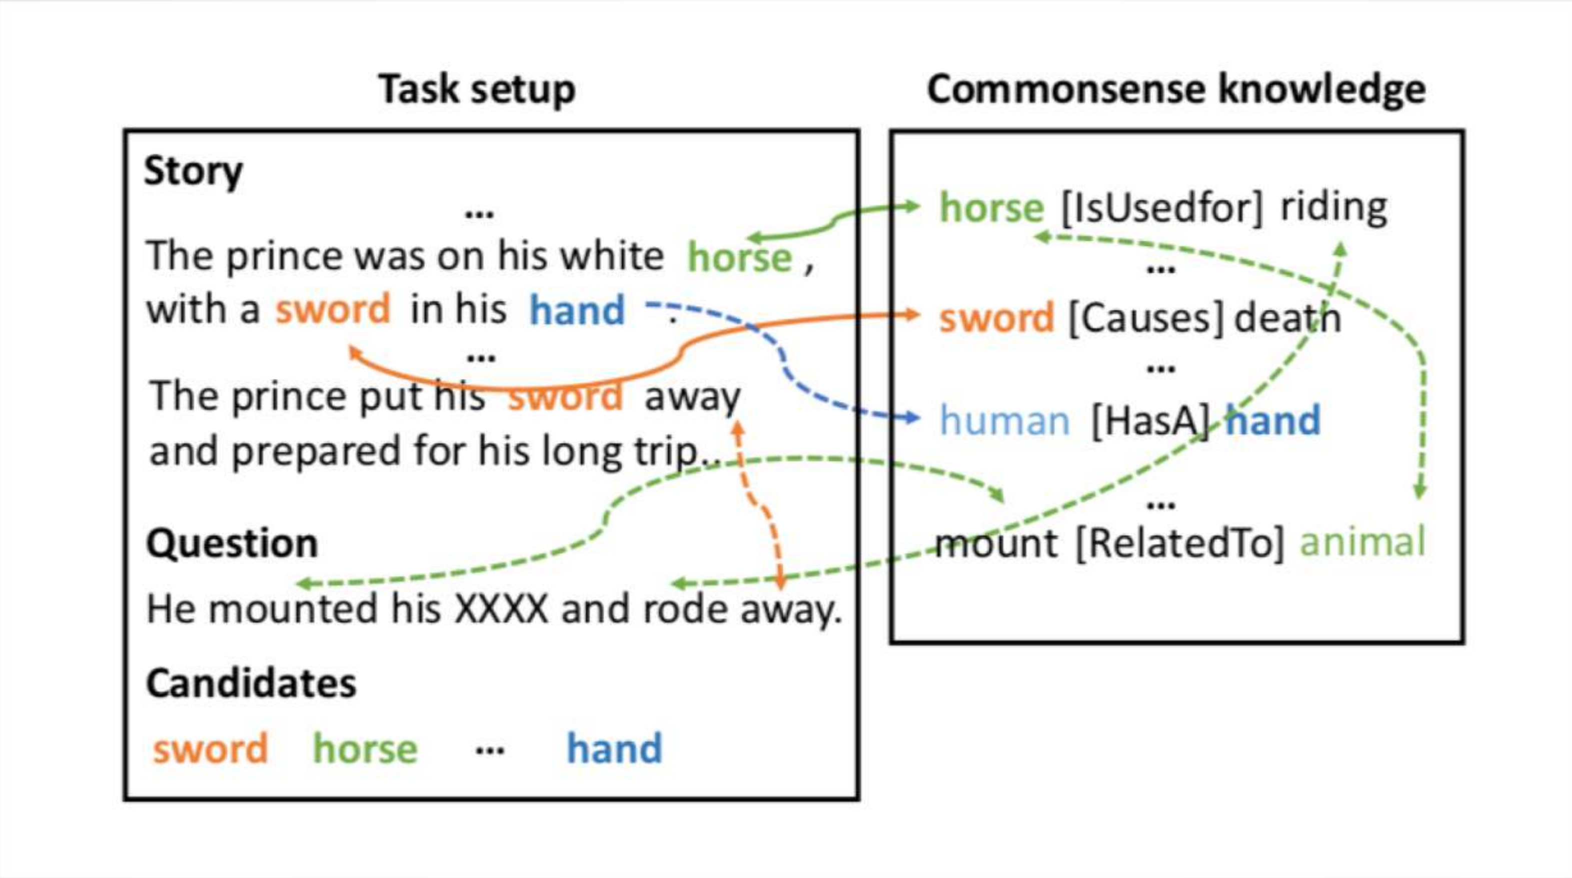

Mihaylov and Frank. 2018. Knowledgeable Reader: Enhancing Cloze-Style Reading Comprehension with External Commonsense Knowledge

CS Know. to Intent, Reaction, Emotion, etc

Event, Intents, and Reactions

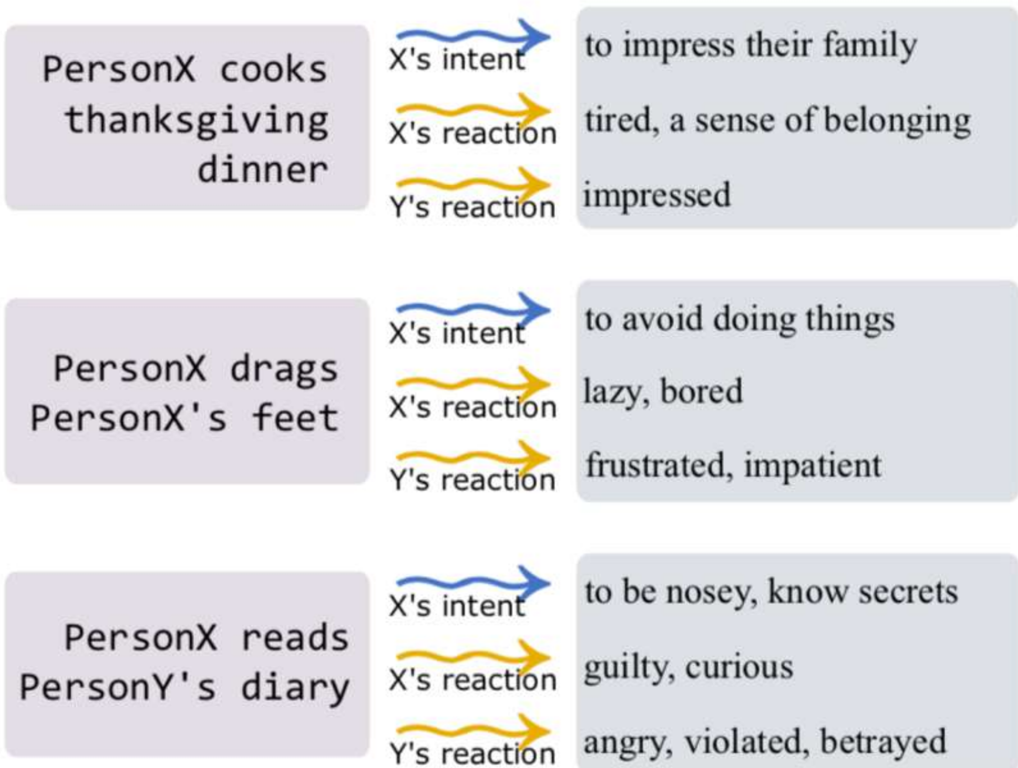

Rashkin et al. 2018. Event2Mind: Commonsense Inference on Events, Intents, and Reactions

Mental states: *motivations* and *emotional reactions*

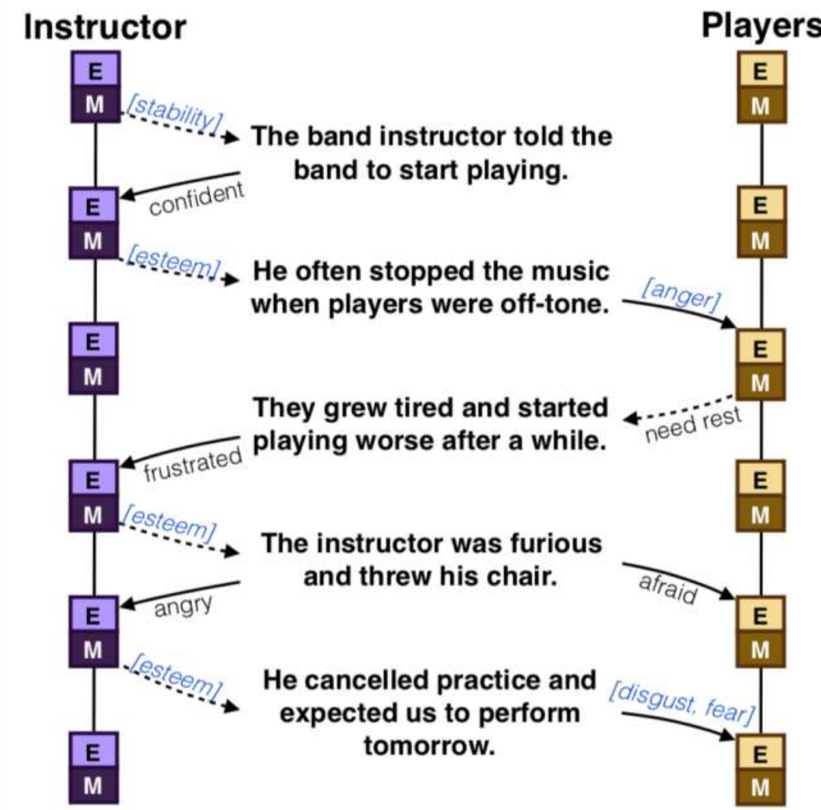

Rashkin et al. 2018. Modeling Naive Psychology of Characters in Simple Commonsense Stories

CS Know. in Language Generation

Dialogue Generation: knowledge

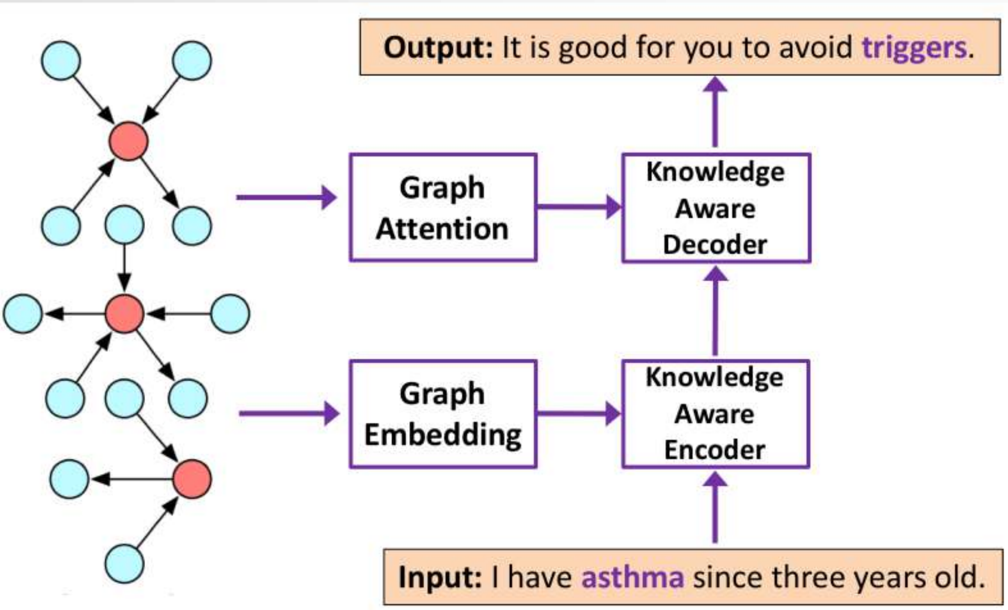

Zhou et al. 2018. Commonsense Knowledge Aware Conversation Generation with Graph Attention.

Story Ending Generation: logic

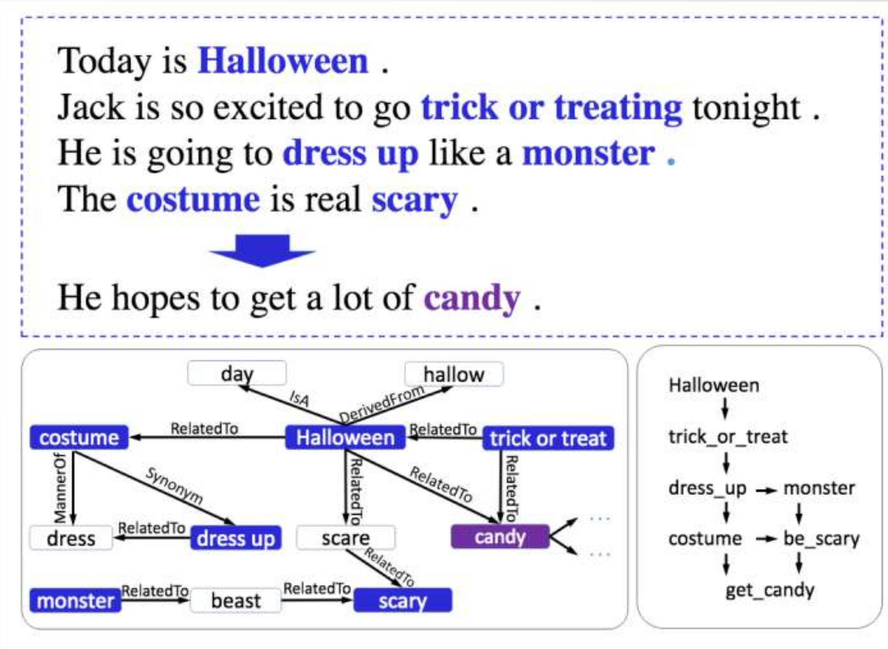

Guan et al. 2019. Story Ending Generation with Incremental Encoding and Commonsense Knowledge.

Commonsense Knowledge

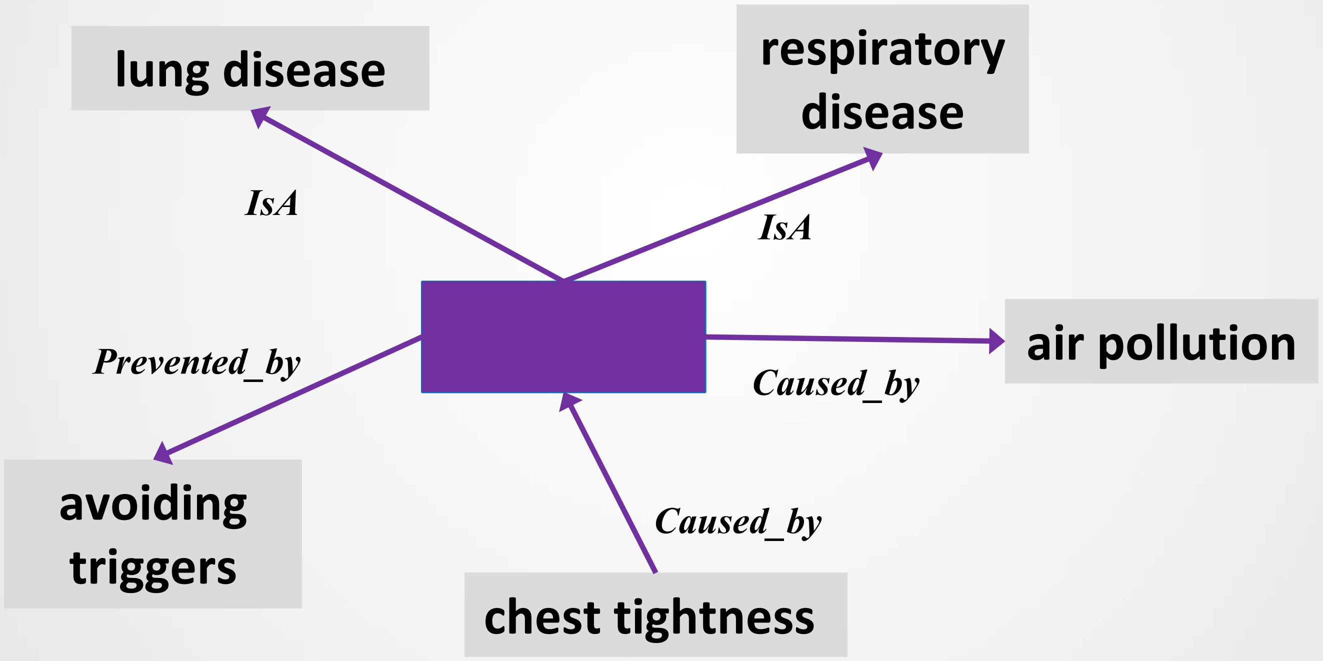

From ConceptNet

Commonsense Knowledge

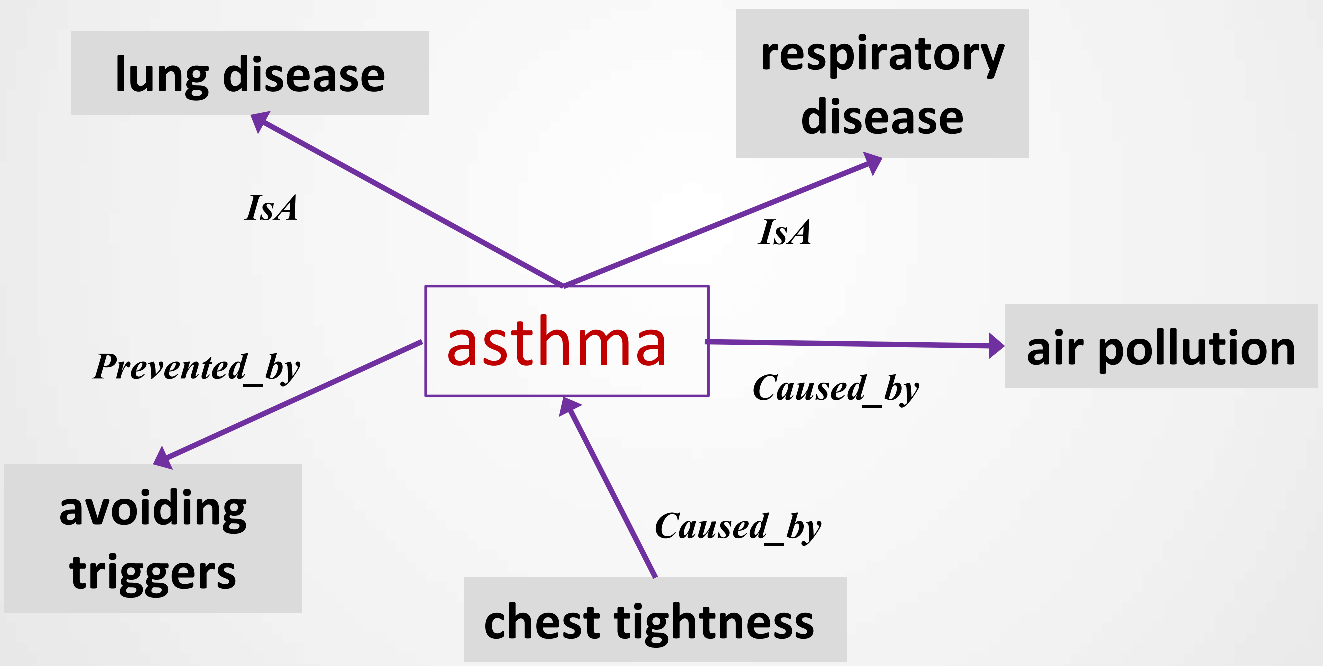

From ConceptNet

Commonsense-aware Dialog Generation

Input: I have an asthma since three years old.

Triples in knowledge graph:
(lung disease, IsA, asthma)
(asthma, Prevented_by, avoiding triggers)

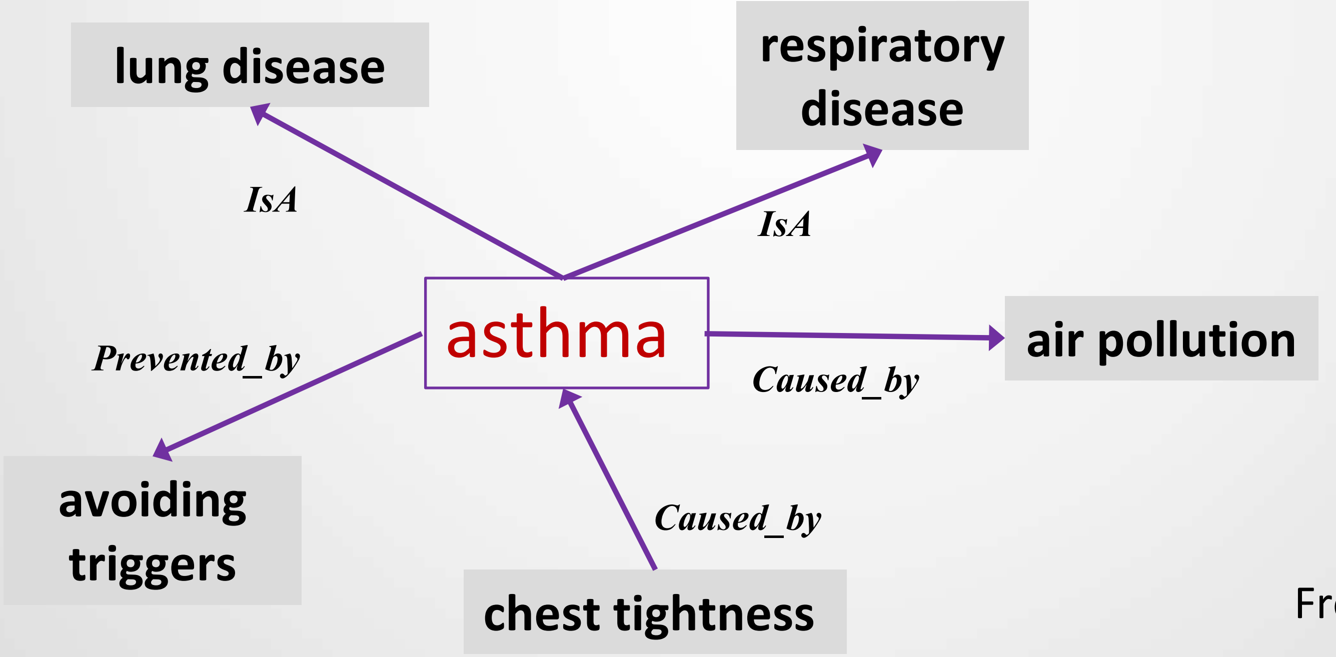

From ConceptNet

Commonsense-aware Dialog Generation

Input: I have an asthma since three years old.

Output: I am sorry to hear that. Maybe avoiding triggers can prevent asthma attacks.

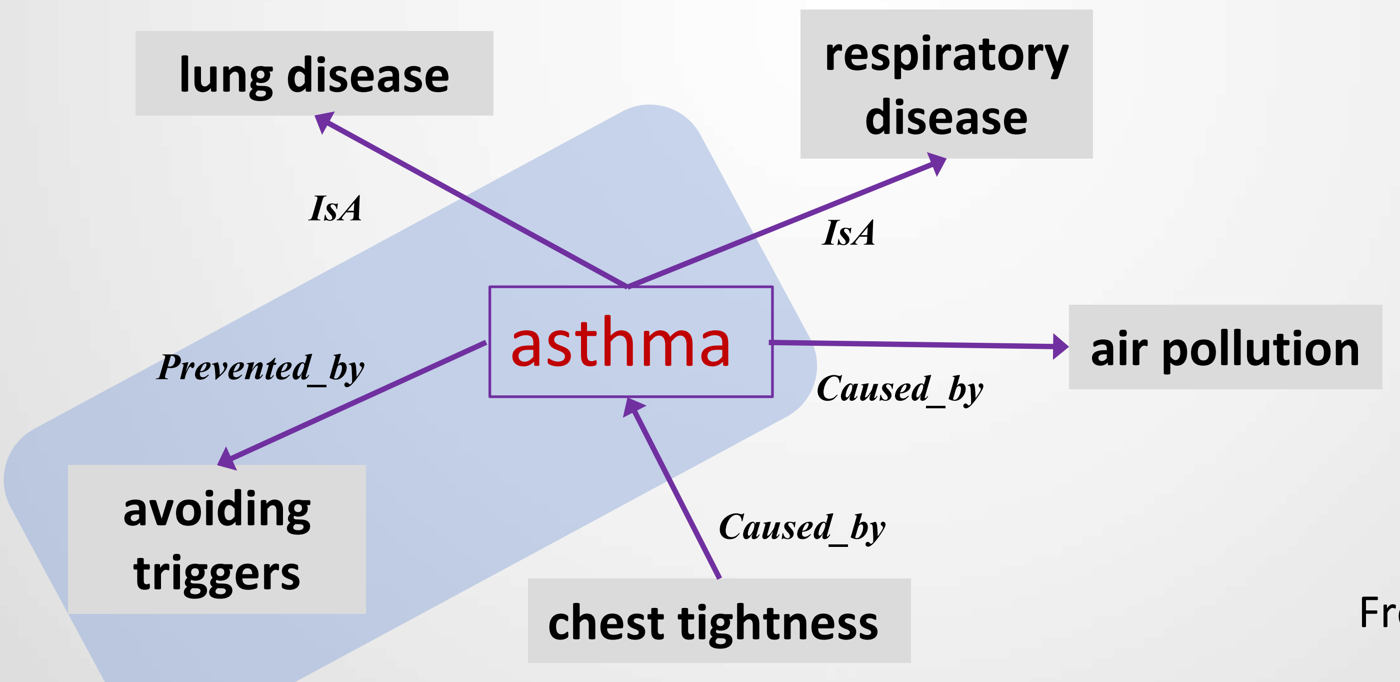

From ConceptNet

Commonsense-aware Dialog Generation

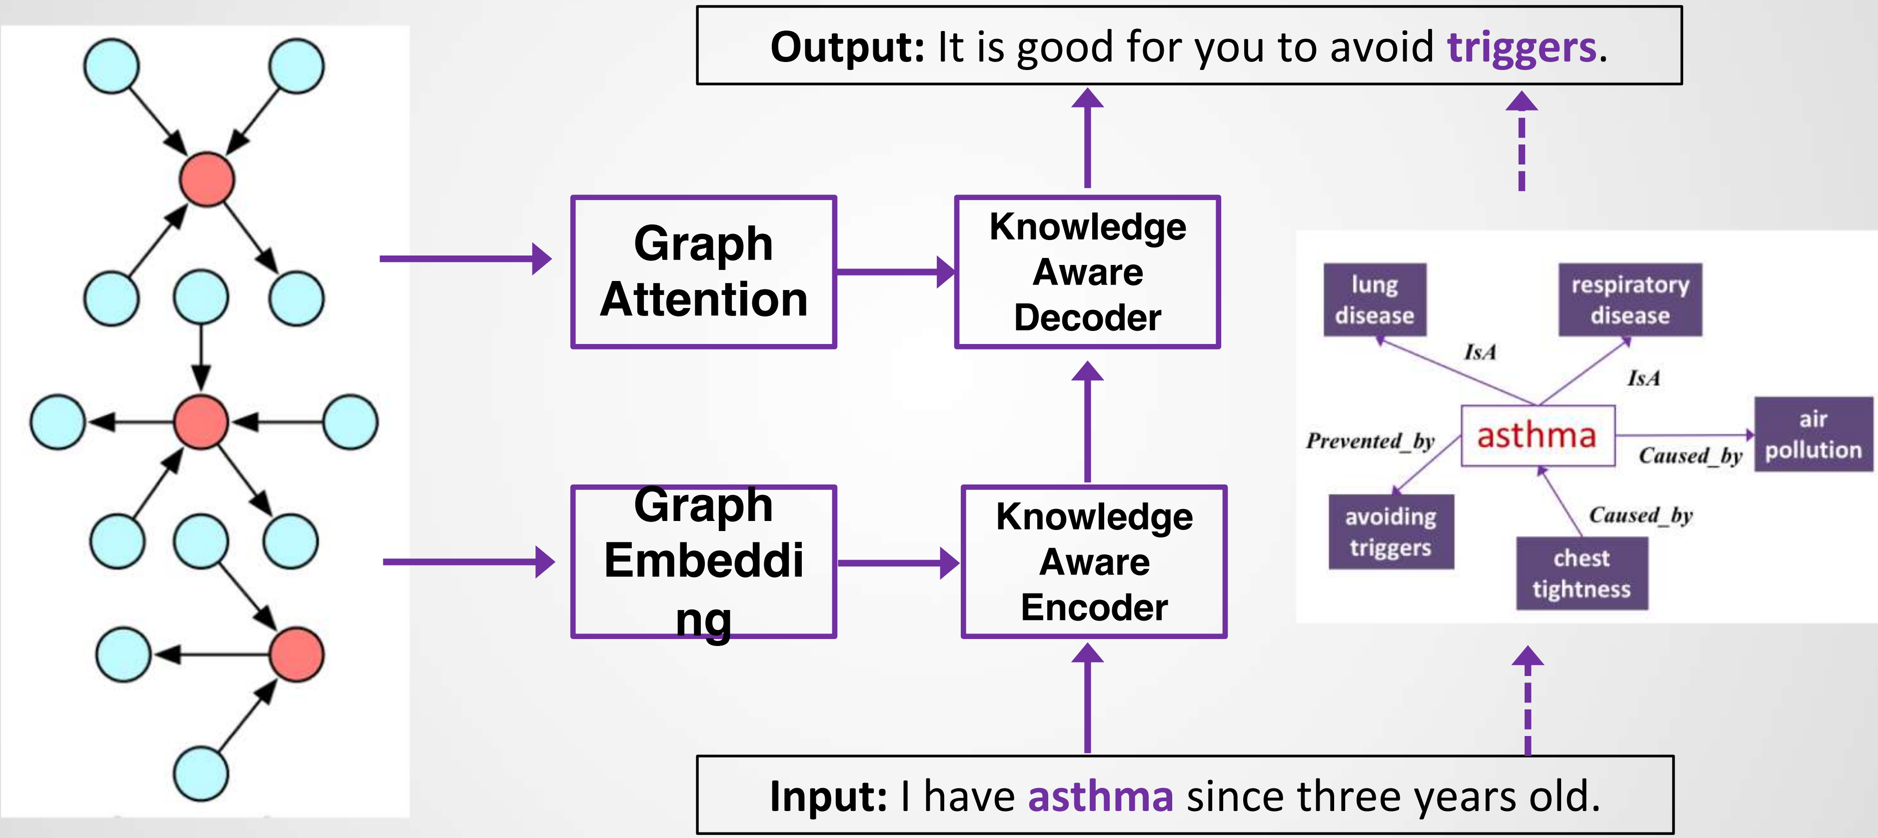

Hao Zhou, Tom Yang, Minlie Huang, Haizhou Zhao, Jingfang Xu, Xiaoyan Zhu. Commonsense Knowledge Aware Conversation Generation with Graph Attention. IJCAI-ECAI 2018, Stockholm, Sweden. **Distinguished paper**

Commonsense-aware Dialog Generation

Encoding with static graph attention: encoding semantics in graph, Feeding knowledge-enhanced info. into the encoder

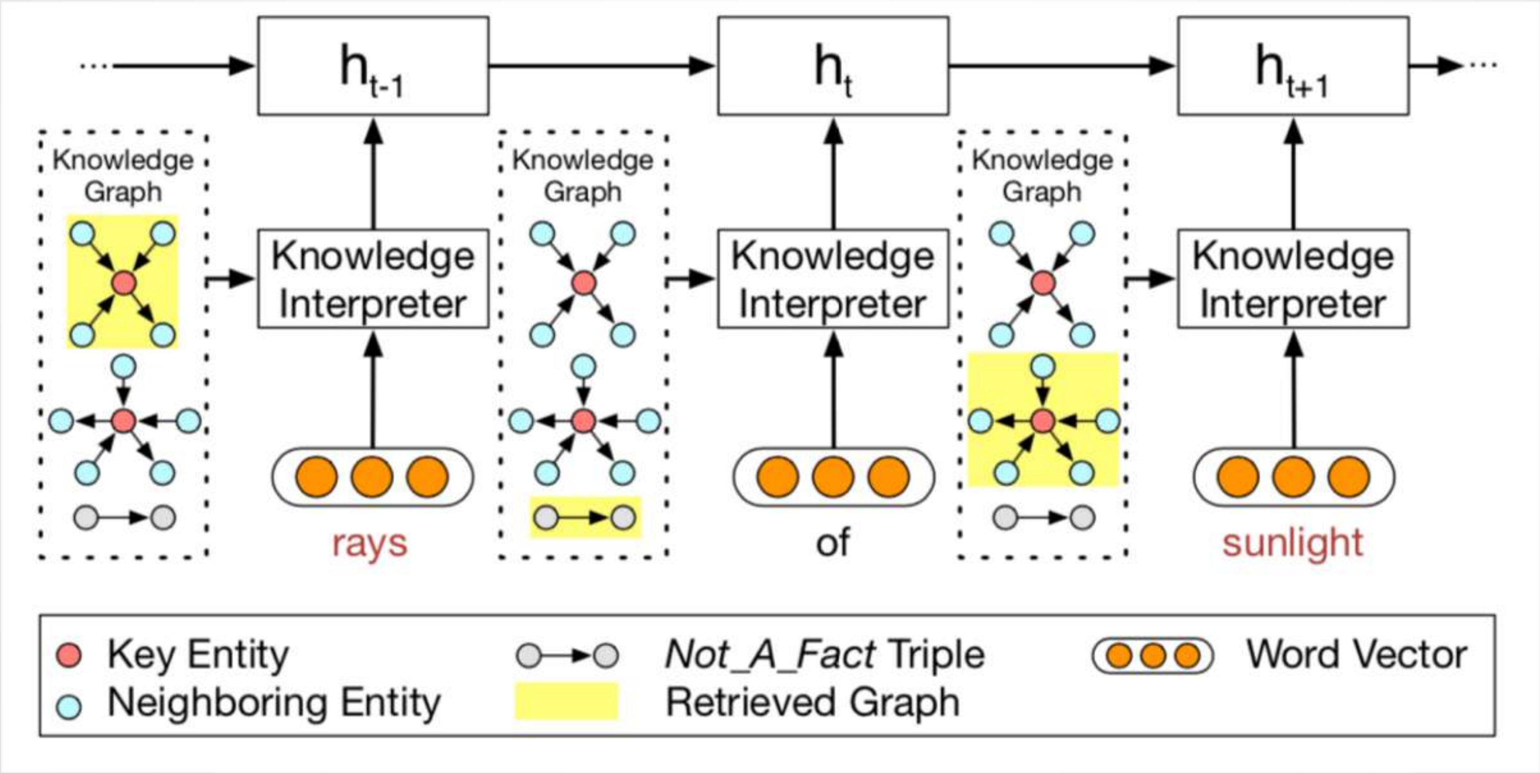

Commonsense-aware Dialog Generation

Decoding with dynamic graph attention: first attend a graph, then to a triple within that graph, finally generate with the words in a graph

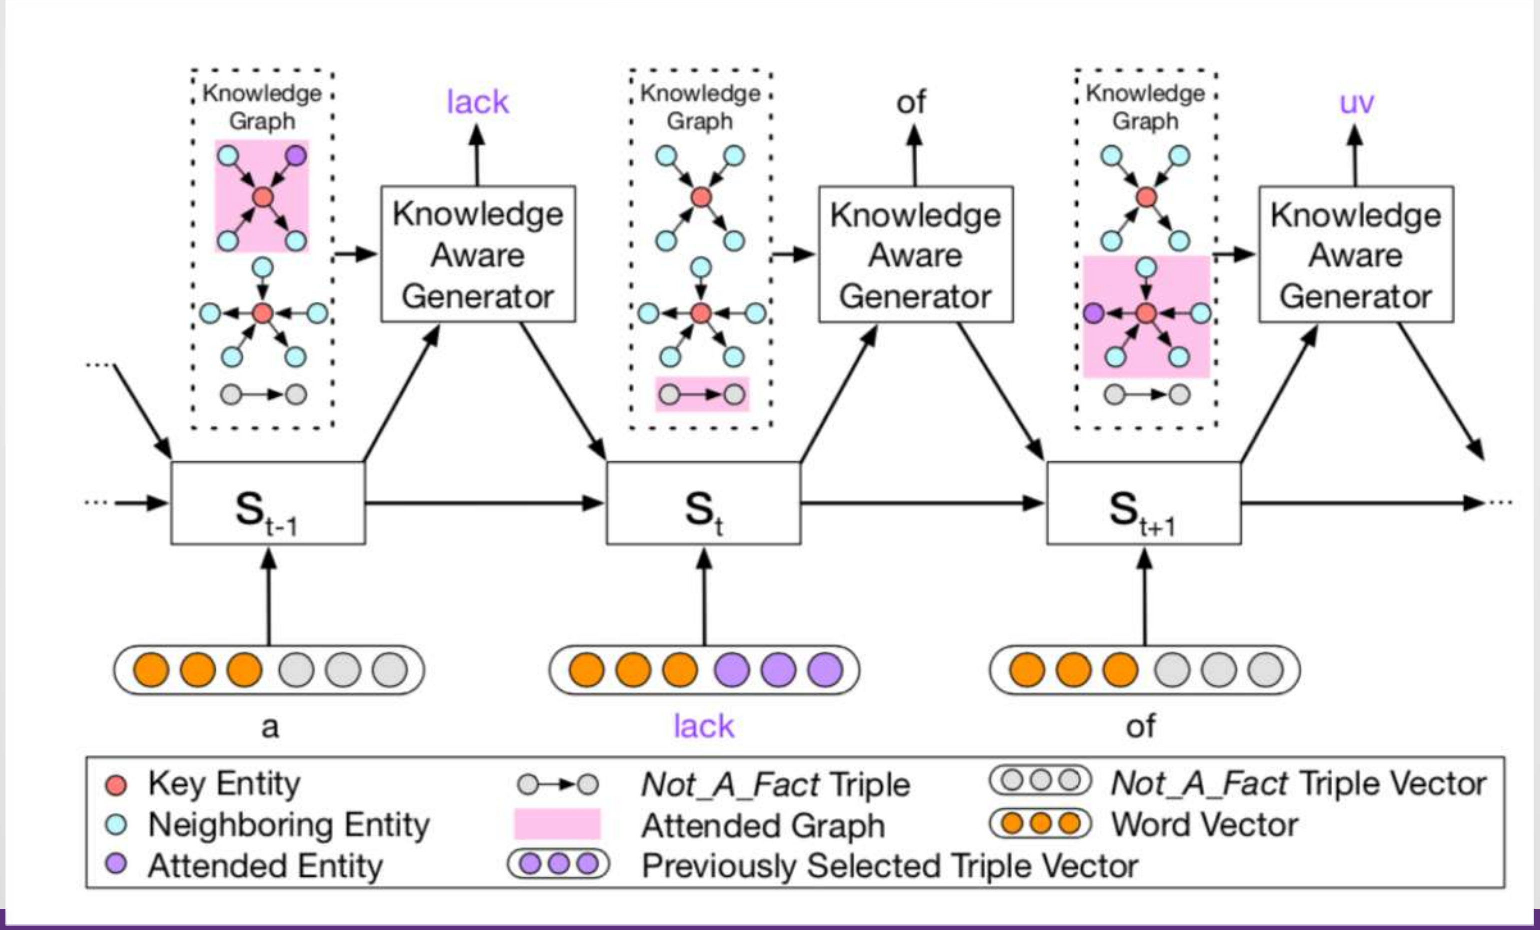

Commonsense-aware Dialog Generation

- Dataset: filtered from 10M reddit single-round dialogs

Conversational Pairs		Commonsense KB	
Training	3,384,185	Entity	21,471
Validation	10,000	Relation	44
Test	20,000	Triple	120,850

Table 1: Statistics of the dataset and the knowledge base.

Commonsense-aware Dialog Generation

Automatic evaluation

Model	Overall		High Freq.		Medium Freq.		Low Freq.		OOV	
	ppx.	ent.	ppx.	ent.	ppx.	ent.	ppx.	ent.	ppx.	ent.
Seq2Seq	47.02	0.717	42.41	0.713	47.25	0.740	48.61	0.721	49.96	0.669
MemNet	46.85	0.761	41.93	0.764	47.32	0.788	48.86	0.760	49.52	0.706
CopyNet	40.27	0.96	36.26	0.91	40.99	0.97	42.09	0.96	42.24	0.96
CCM	**39.18**	**1.180**	**35.36**	**1.156**	**39.64**	**1.191**	**40.67**	**1.196**	**40.87**	**1.162**

Manual evaluation (Sign-test, p-value<0.005)

Model	Overall		High Freq.		Medium Freq.		Low Freq.		OOV	
	app.	inf.	app.	inf.	app.	inf.	app.	inf.	app.	inf.
CCM vs. Seq2Seq	0.616	0.662	0.605	0.656	0.549	0.624	0.636	0.650	0.673	0.716
CCM vs. MemNet	0.602	0.647	0.593	0.656	0.566	0.640	0.622	0.635	0.626	0.657
CCM vs. CopyNet	0.600	0.640	0.606	0.669	0.586	0.619	0.610	0.633	0.596	0.640

Generated Example

Post: He proposed March 5th. We will be married October 10th. So 7 months

Knowledge: (engaged, RelatedTo, married), (wedding, RelatedTo, married), (miss, DistinctFrom, married), (March, RelatedTo, months)

- **Seq2Seq**: I'm not sure how to answer that. I'm not sure how to make a date.
- **MemNet**: I'm assuming he's married to a gay marriage?
- **CopyNet**: I'm sorry, I'm not sure what you mean by that.
- **CCM**: I'm pretty sure he's going to be **engaged** in the next 6 **months**.

Generated Example

Post: Totally thought it was going to be doug's grave.
Knowledge: (thought, RelatedTo, thinking), (doug, HasProperty, dead), (dead, RelatedTo, grave), (grave, RelatedTo, death)

- **Seq2Seq**: I was thinking the same thing.
- **MemNet**: I thought it was a dead end.
- **CopyNet**: I think it's a bit of a stretch.
- **CCM**: I thought it was doug's **death**. I was **thinking** the same thing.

Story Ending Generation with CS Know.

Finding context clues: the temporal order of events and entities.

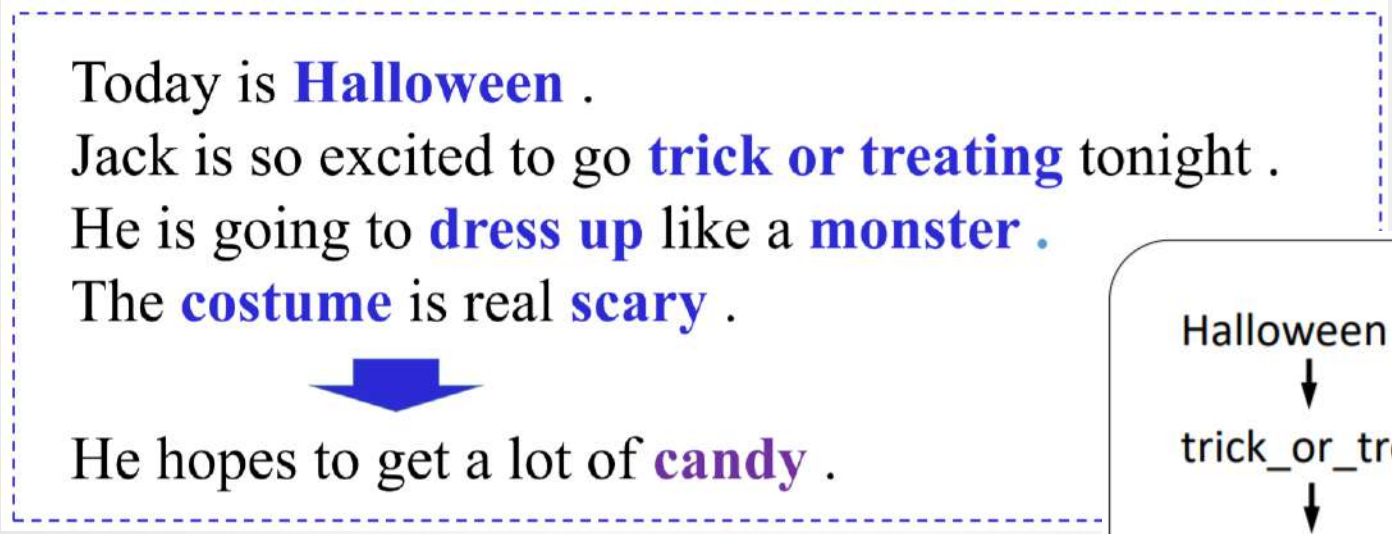

Jian Guan, Yansen Wang, Minlie Huang. Story Ending Generation with Incremental Encoding and Commonsense Knowledge. AAAI 2019

Story Ending Generation with CS Know.

Knowing the causality with **commonsense knowledge**

Today is **Halloween** .
Jack is so excited to go **trick or treating** tonight .
He is going to **dress up** like a **monster** .
The **costume** is real **scary** .

He hopes to get a lot of **candy** .

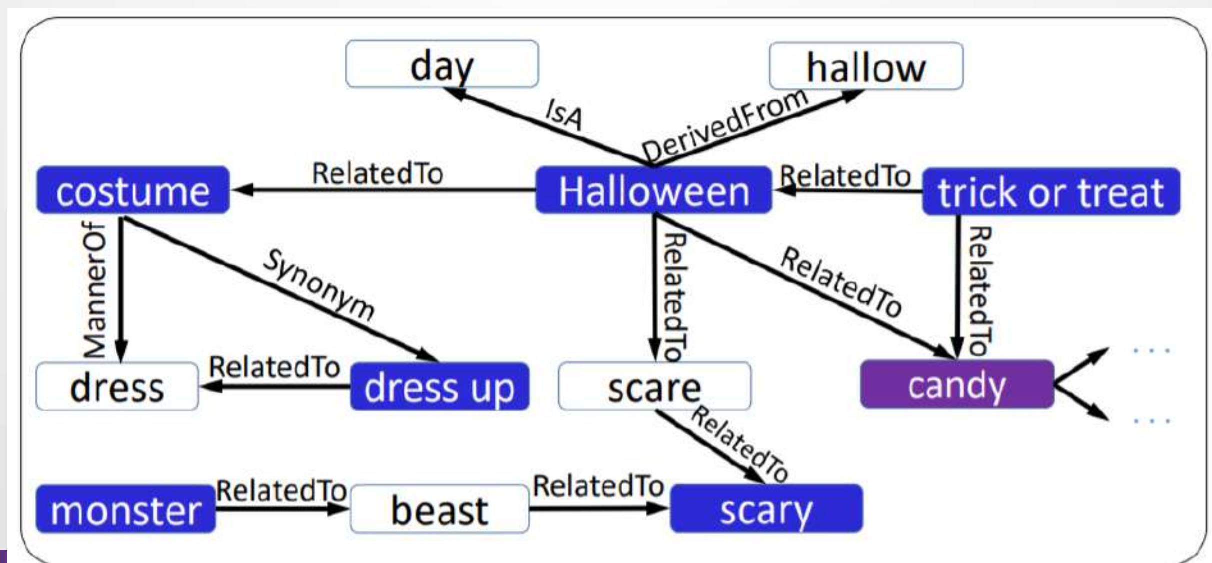

Story Ending Generation with CS Know.

Incremental Encoding

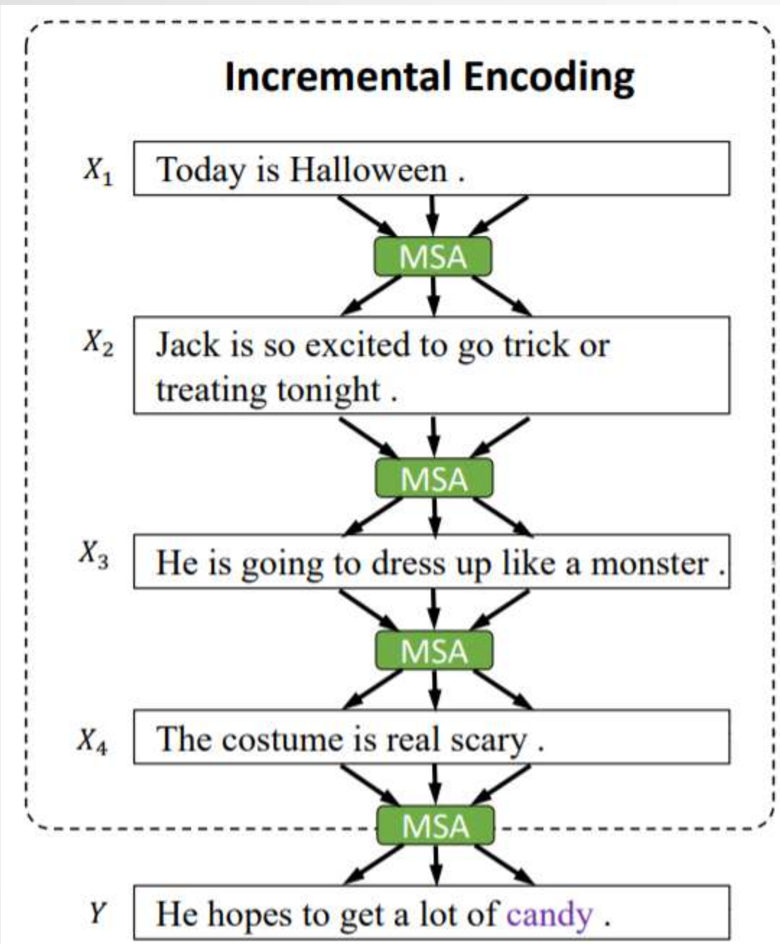

Multi-Source Attention

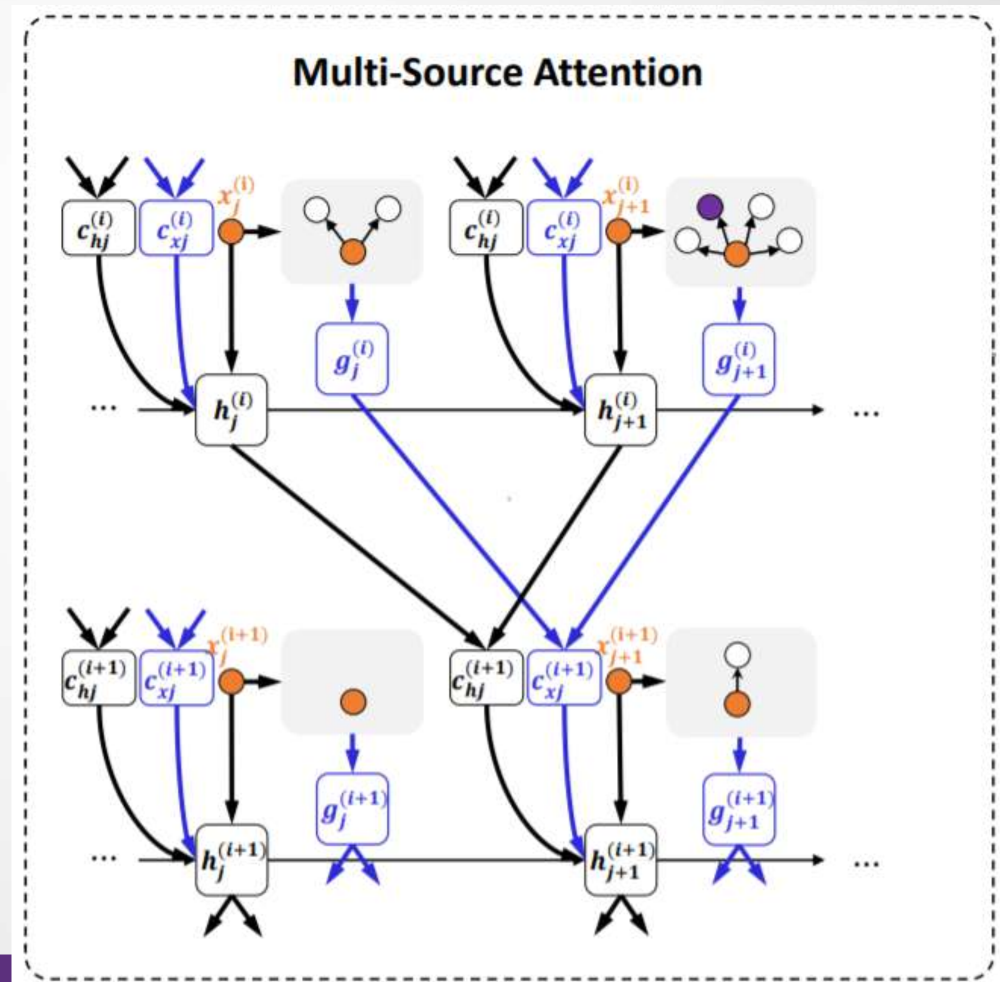

Story Ending Generation with CS Know.

Attention to the knowledge base: static graph attention

Graph attention

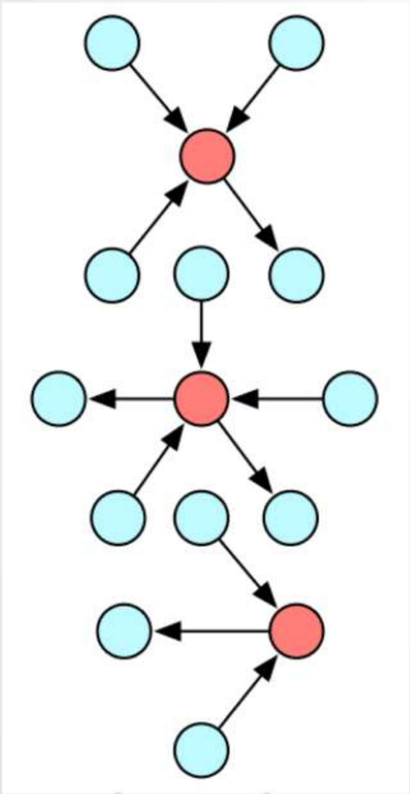

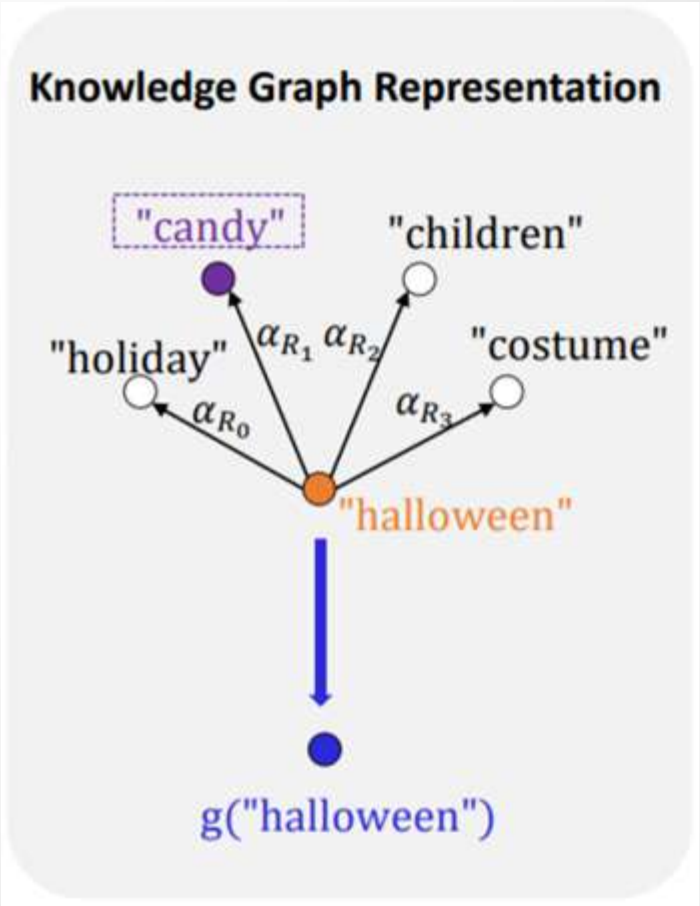

Story Ending Generation with CS Know.

- **Graph Attention**

$$\mathbf{g}(x) = \sum_{i=1}^{N_x} \alpha_{R_i} [\mathbf{h}_i; \mathbf{t}_i],$$

$$\alpha_{R_i} = \frac{e^{\beta_{R_i}}}{\sum_{j=1}^{N_x} e^{\beta_{R_j}}},$$

$$\beta_{R_i} = (\mathbf{W_r r}_i)^{\mathrm{T}} \, tanh(\mathbf{W_h h}_i + \mathbf{W_t t}_i),$$

- **Contextual Attention**

$$\mathbf{g}(x) = \sum_{i=1}^{N_x} \alpha_{R_i} \mathbf{M}_{R_i},$$

$$\mathbf{M}_{R_i} = BiGRU(\mathbf{h}_i, \mathbf{r}_i, \mathbf{t}_i),$$

$$\alpha_{R_i} = \frac{e^{\beta_{R_i}}}{\sum_{j=1}^{N_x} e^{\beta_{R_j}}},$$

$$\beta_{R_i} = \mathbf{h}_{(x)}^{\mathrm{T}} \mathbf{W_c M}_{R_i},$$

Generated Examples

Story 1:

Context:

Taj has never drank an espresso drink.
He ordered one while out with his friends.
The shot of espresso tasted terrible to him.
Taj found that he couldn't stop talking or moving.

Generated Ending:

He decided to never drink again.

Story 2:

Context:

Martha is cooking a special meal for her family.
She wants everything to be just right for when they eat.
Martha perfects everything and puts her dinner into the oven.
Martha goes to lay down for a quick nap.

Generated Ending:

When she gets back to the kitchen, she sees a burning light on the stove.

An Example of "Logic Chains"

Building context clues incrementally

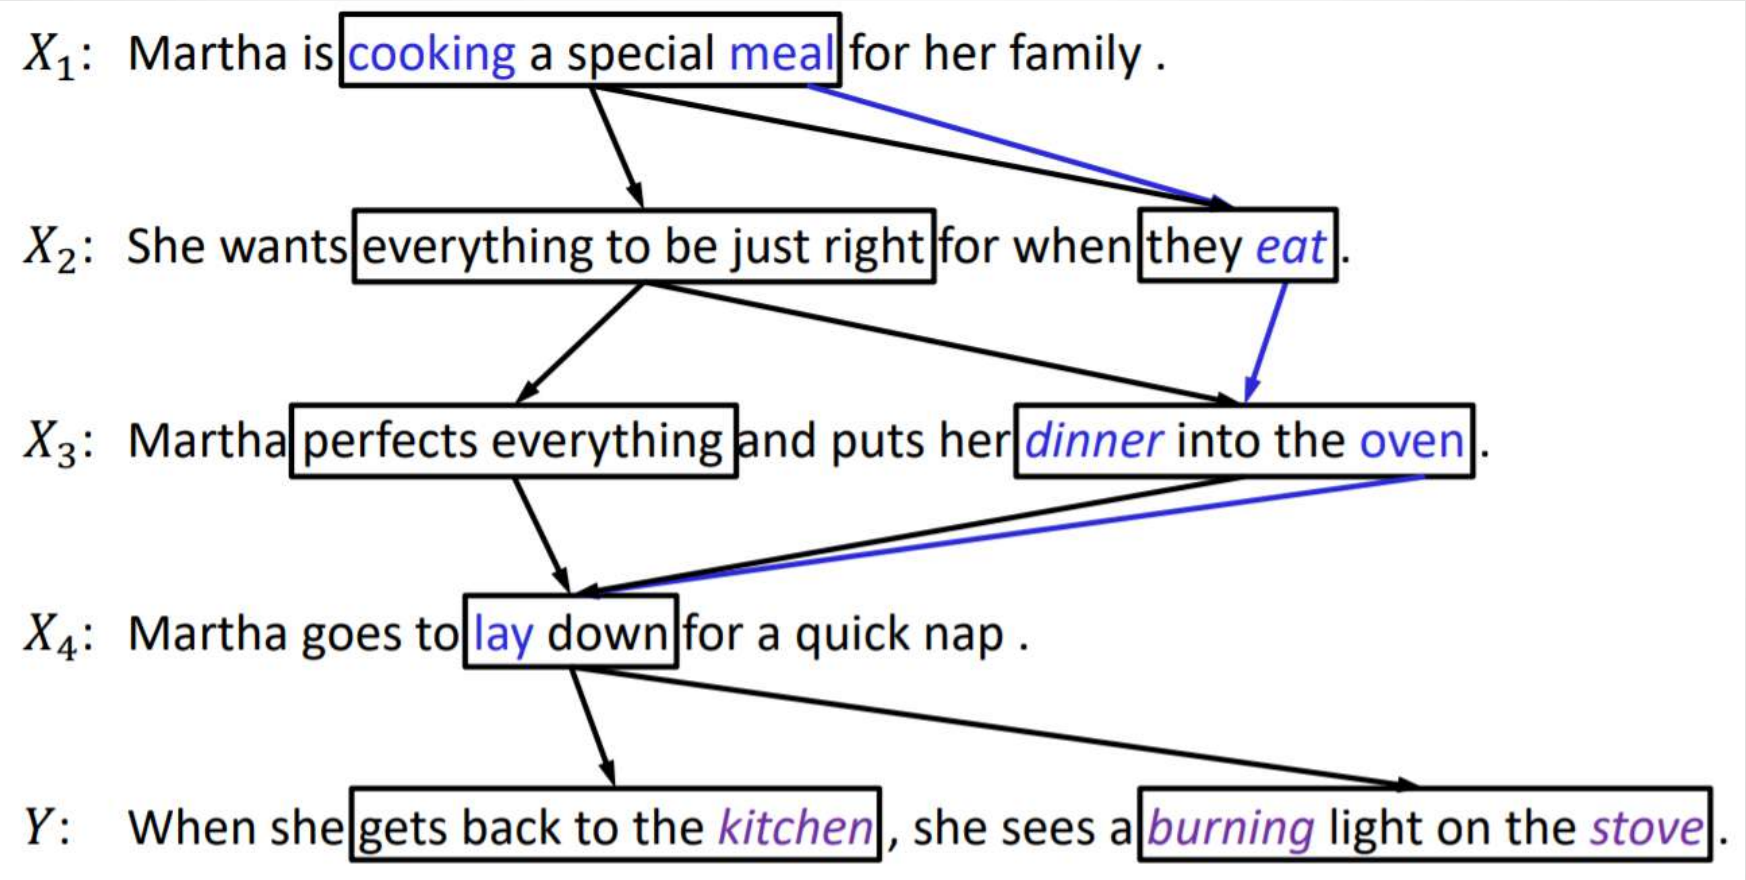

Controllable Language Generation

- Three **fundamental problems** in current neural language generation models
  - **Semantics** (real understanding)
  - **Consistency** (long text generation)
  - **Logic** (reasonable and making sense)

- New architecture: **symbolic knowledge** + **planning** + **neural computing**

Future: Knowledge-grounded Social Chatbot

Name	The Shape of Water
Year	2017
Director	Guillermo del Toro
Genre	Fantasy, Drama
Cast	Sally Hawkins as Elisa Esposito, a mute cleaner who works at a secret government laboratory. Michael Shannon as Colonel Richard Strickland, a corrupt military official, Richard Jenkins as Giles, Elisa's closeted neighbor and close friend who is a struggling advertising illustrator. Octavia Spencer as Zelda Delilah Fuller, Elisa's co-worker and friend who serves as her interpreter., Michael Stuhlbarg as Dimitri Mosenkov, a Soviet spy working as a scientist studying the creature, under the alias Dr. Robert Hoffstetler.

A Dataset for Document Grounded Conversations (Zhou et al., 2018)

Future: Knowledge-grounded Social Chatbot

User 2:	I thought The Shape of Water was one of Del Toro's best works. What about you?
User 1:	Did you like the movie?
User 1:	Yes, his style really extended the story.
User 2:	I agree. He has a way with fantasy elements that really helped this story be truly beautiful.

Table 2: An example conversation for scenario 2. Both User 1 and User 2 have access to the Wiki document. The full dialogue is attached in the Appendix.

A Dataset for Document Grounded Conversations (Zhou et al., 2018)

Future: Knowledge-grounded Social Chatbot

Topic:	Lifeguard
Apprentice:	So I am a lifeguard. Know anything about saving lives in water?
Wizard:	I'm impressed! It's a big responsibility to supervise other people's safety in the water! Tell me more.
Apprentice:	Well, I help make sure people do not drown or get injured while in or near the water!
Knowledge:	A lifeguard is a rescuer who supervises the safety and rescue of swimmers, surfers, ... Lifeguards are strong swimmers and trained in CPR/AED first aid, certified in water ... ... In some areas, the lifeguard service also carries out mountain rescues, or may function as the primary EMS provider.
Wizard:	I've heard that in some places, lifeguards also help with other sorts of emergencies, like mountain rescues! Is that part of your job too?
Apprentice:	I have! I feel like you know much about this! What brings you to know so much?
Wizard:	Oh, that's about the extent of my knowledge. I've just been around beaches and I've always admired lifeguards. I'm not a super strong swimmer myself.

Wizard of Wikipedia: Knowledge-Powered Conversational agents. Dinan et al. 2018.

感谢关注！

aihuang@tsinghua.edu.cn
Http://coai.cs.tsinghua.edu.cn/hml