Neurosymbolic Knowledge Bases

Lara J. Martin (she/they)

https://laramartin.net/interactive-fiction-class

Modified from slides from the [ACL 2020 Commonsense Tutorial] by Yejin Choi, Vered Shwartz, Maarten Sap, Antoine Bosselut, and Dan Roth

Learning Objectives

Define what neurosymbolic methods are

Follow examples of integrated and post-hoc knowledge base integration

Compare GPT-3-era neurosymbolic systems to modern neural systems

Review: Desirable properties for a commonsense resource

COVERAGE

Large scale

Diverse knowledge types

USEFUL

High quality knowledge

Usable in downstream tasks

Multiple resources tackle different knowledge types

Review:
Eliciting commonsense from humans

EXPERTS CREATE KNOWLEDGE BASE

Advantages:

- Quality guaranteed
- Can use complex representations (e.g., CycL, LISP)

Drawbacks:

- Time cost
- Training users

OpenCyc 4.0 (Lenat, 2012)

WordNet (Miller et al., 1990)

NON-EXPERTS WRITE KNOWLEDGE IN NATURAL LANGUAGE PHRASES

Natural language

- Accessible to non-experts
- Different phrasings allow for more nuanced knowledge

Fast and scalable collection

- Crowdsourcing
- Games with a purpose

ATOMIC (Sap et al., 2019)

ConceptNet 5.5 (Speer et al., 2017)

Neurosymbolic Methods

The combination of neural networks (“**neuro**”) and older, **symbolic** AI methods

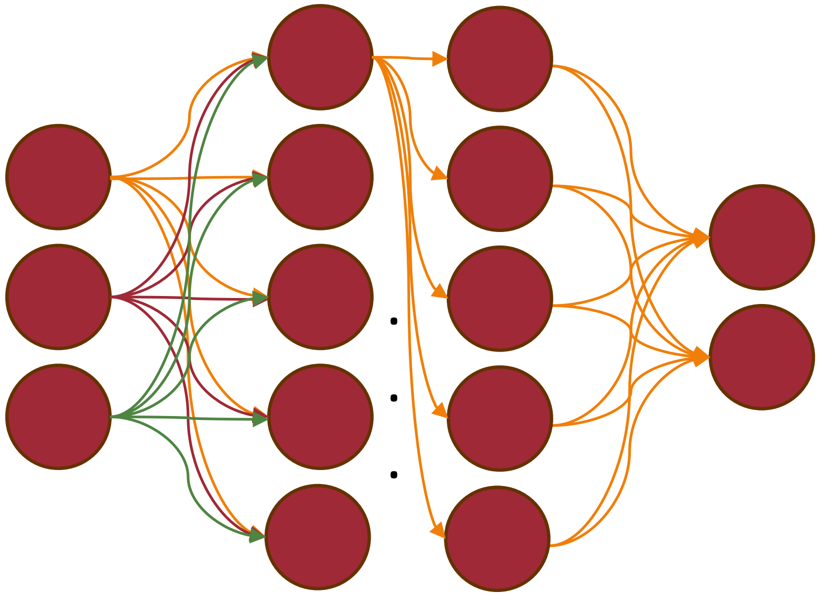

+

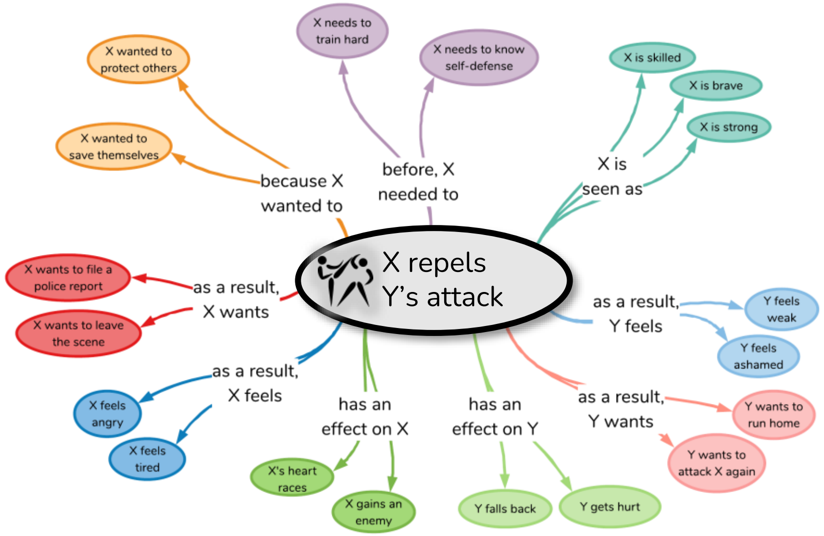

M. Sap *et al.*, “ATOMIC: An Atlas of Machine Commonsense for If-Then Reasoning,” *AAAI Conference on Artificial Intelligence (AAAI)*, vol. 33, no. 1, pp. 3027–3035, 2019, doi: 10.1609/aaai.v33i01.33013027.

Why combine them?

NEURAL NETWORKS

Statistical patterns over data

Easy to generate new text from

Need a lot of data to train (and might need to be labeled)

Hard to control

Examples: sequence-to-sequence networks, transformers (LLMs)

SYMBOLIC METHODS

Structured information

Easy for people to understand (interpretable)

Hard to make

- Need experts or a lot of time

Limited set of information

Examples: knowledge bases, planning domains/problems, scripts

Ways of combining them

During training
- Such as in reinforcement learning or retrieval-augmented generation (RAG)

After training
- Like a symbolic “wrapper” – helps validate what the NN is doing

Others??

Ways of combining them

During training

- Such as in reinforcement learning or retrieval-augmented generation (RAG)

After training

- Like a symbolic "wrapper" – helps validate what the NN is doing

Others??

Adding neural networks to knowledge bases

Katrina had the financial means to afford a new car while Monica did not, since _____ had a high paying job.

WINOGRANDE: An Adversarial Winograd Schema Challenge at Scale.
Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. AAAI 2020.

Neural Architecture

[CLS] Katrina had the financial means to afford a new car
while Monica did not, since [SEP] Katrina had a high paying job.

[CLS] Katrina had the financial means to afford a new car
while Monica did not, since [SEP] Monica had a high paying job.

0.51
0.49

Masked Language Models

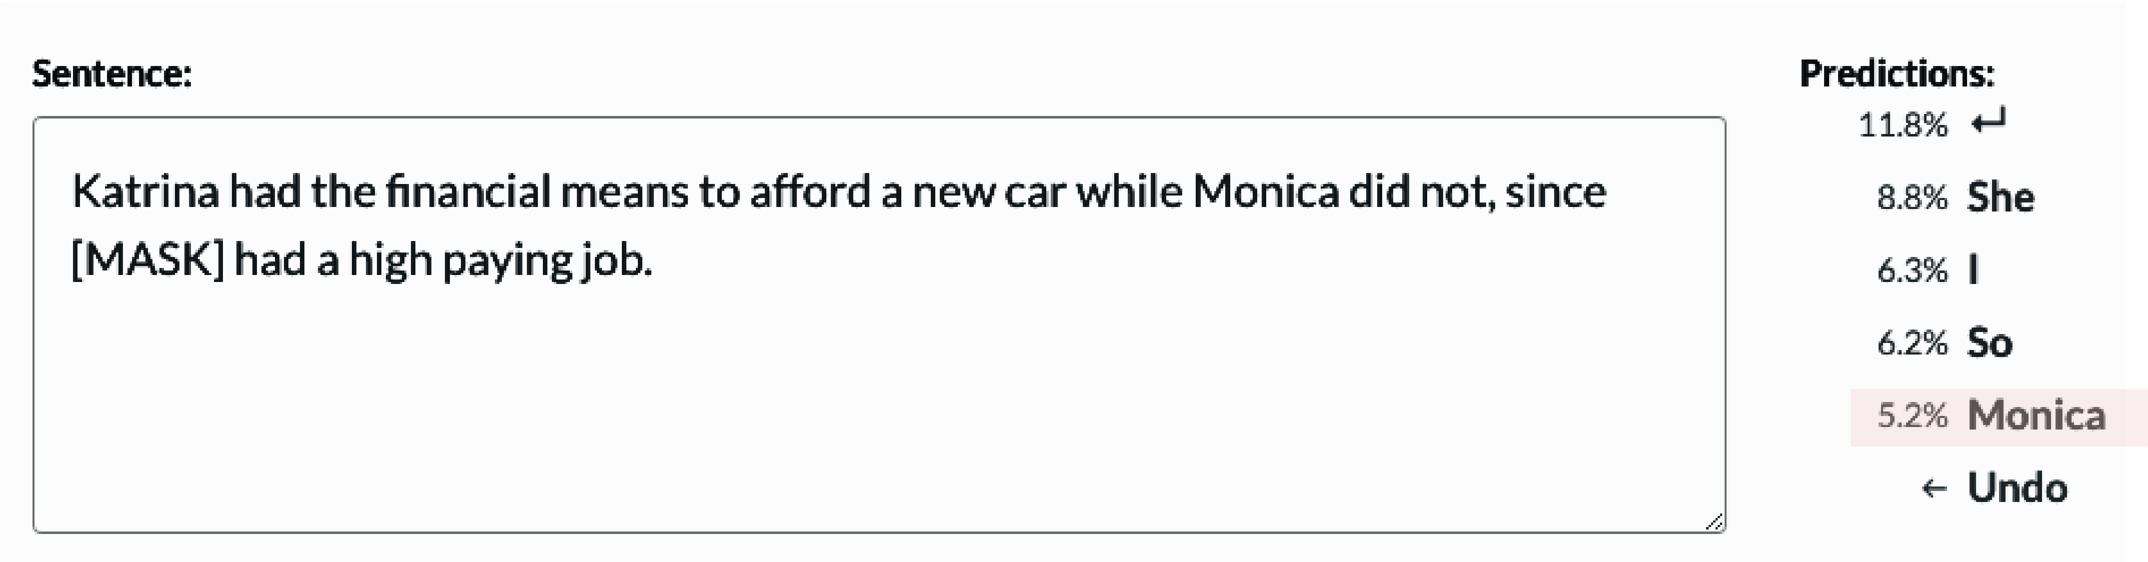

~~https://demo.allennlp.org/masked-lm~~

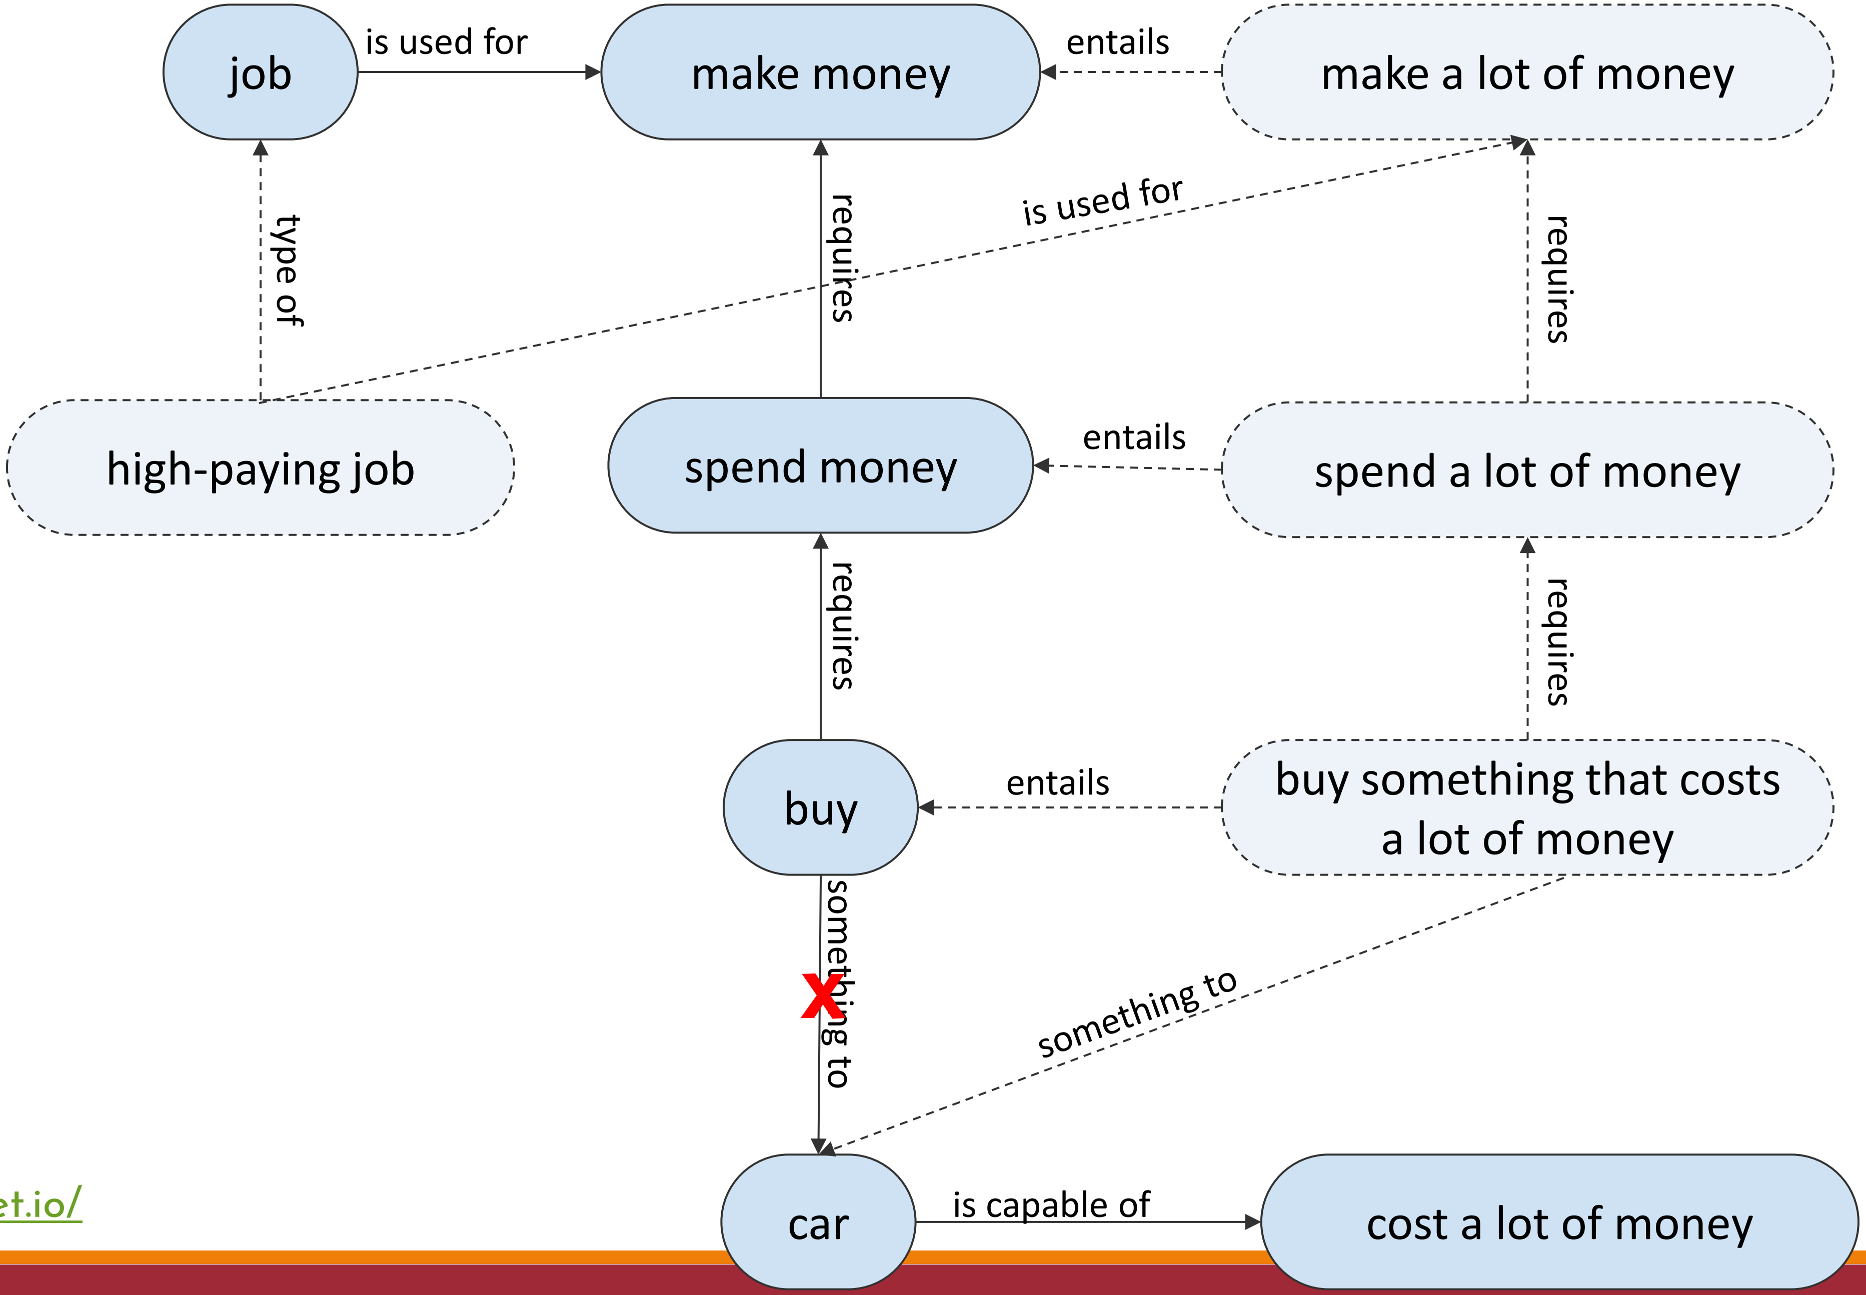

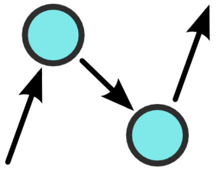

https://conceptnet.io/

Incorporating External Knowledge into Neural Models

General Idea

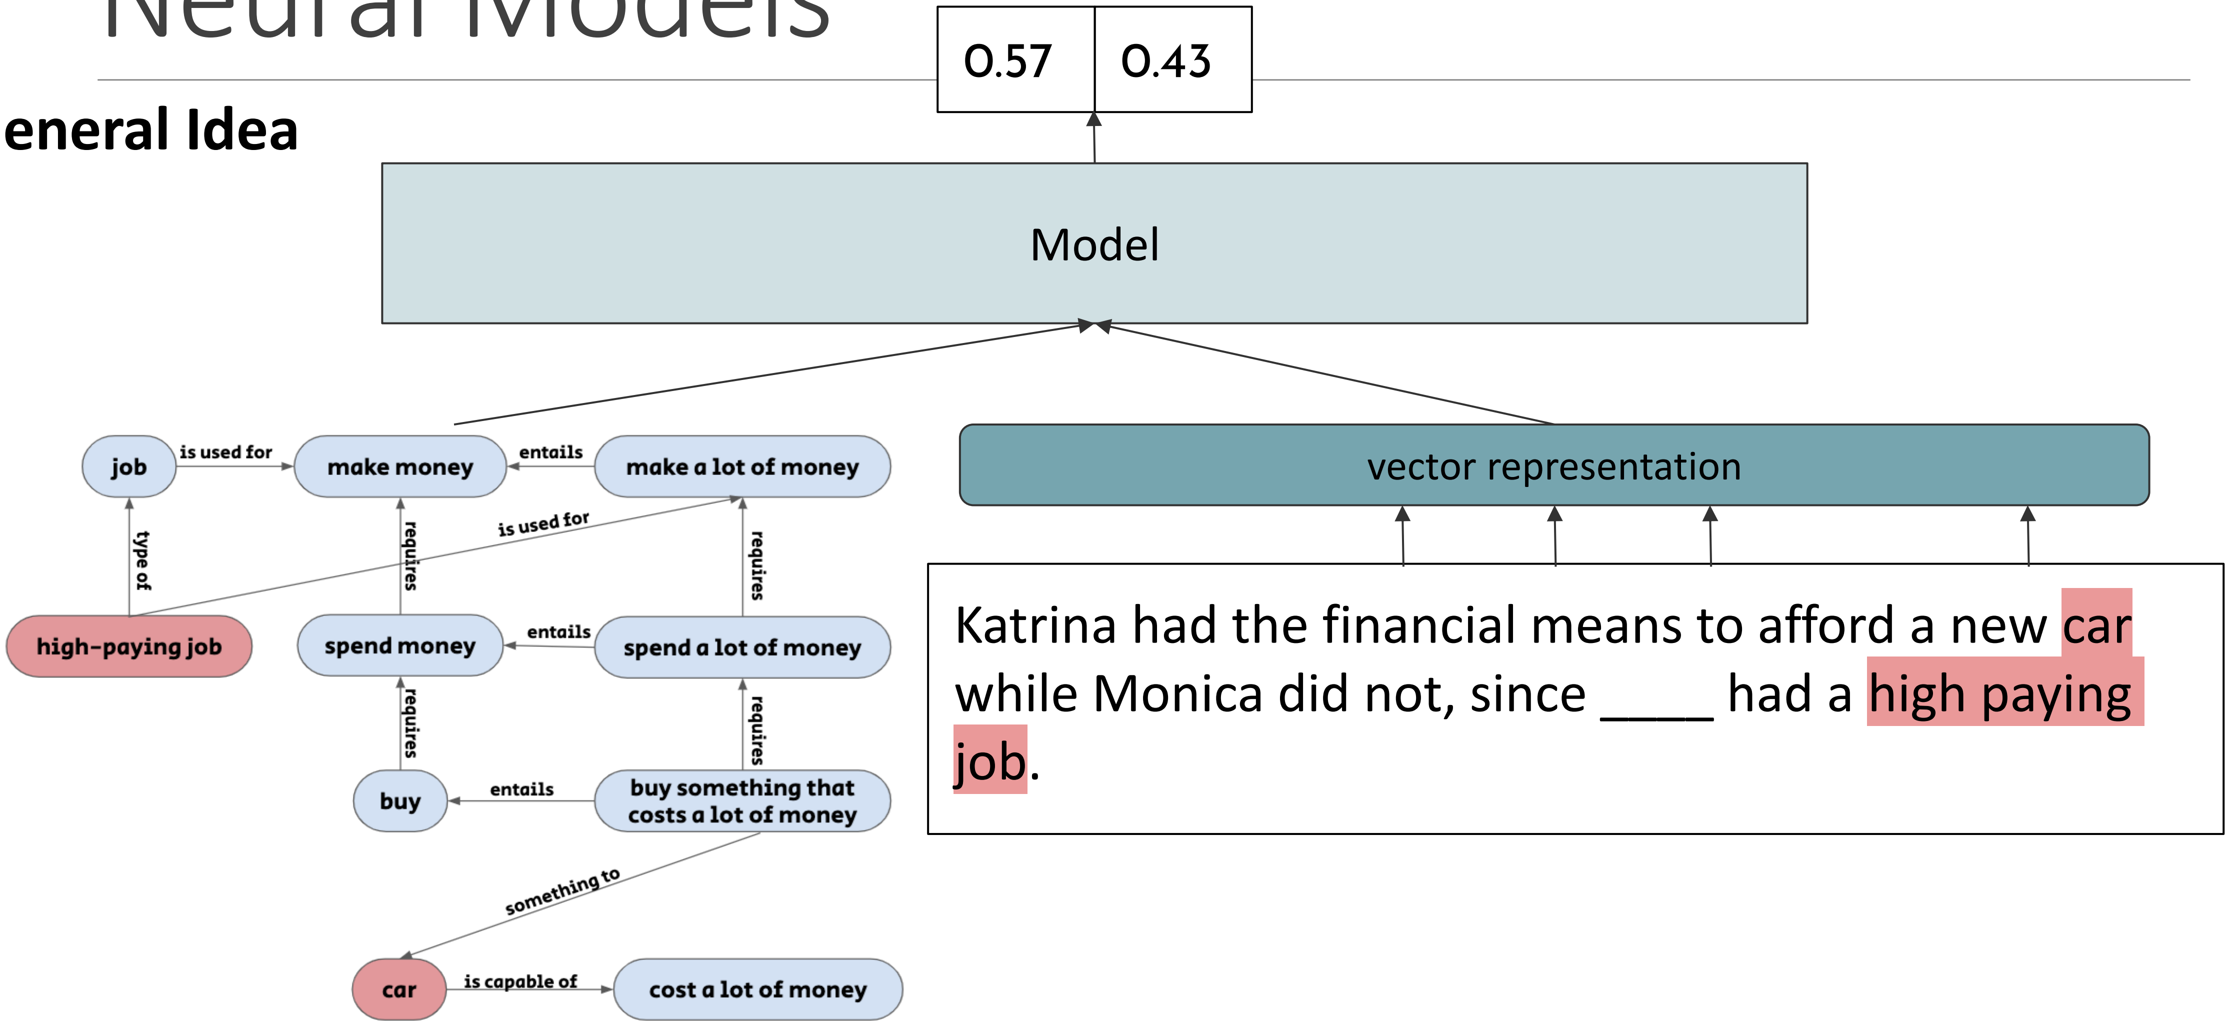

Incorporating External Knowledge into Neural Models

Recipe

Task

Story ending,
Machine Comprehension
Social common sense
NLI

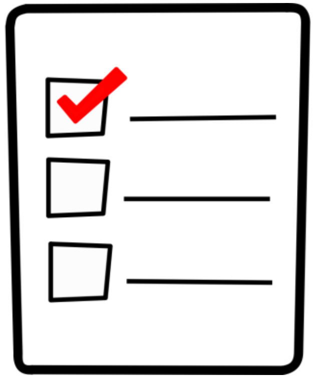

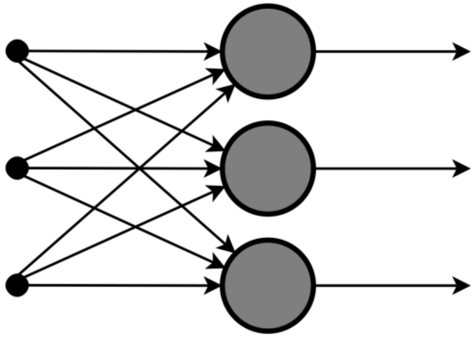

Neural Component

Pre/post pre-trained language models

Knowledge Source

Knowledge bases, extracted from text, hand-crafted rules

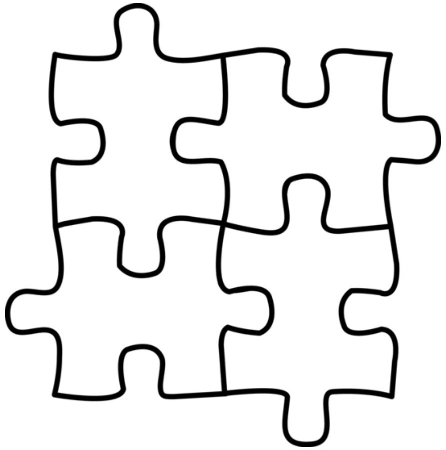

Combination Method

Attention, pruning, word embeddings, multi-task learning

Story Cloze Test

Agatha had always wanted pet birds.
So one day she purchased two pet finches.
Soon she couldn't stand their constant noise.
And even worse was their constant mess.

Agatha decided to buy two more. (Wrong)
Agatha decided to return them. (Right)

Task

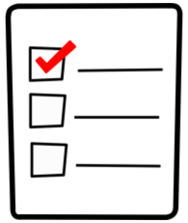

A Corpus and Cloze Evaluation for Deeper Understanding of Commonsense Stories. *Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen.* NAACL 2016.

ConceptNet

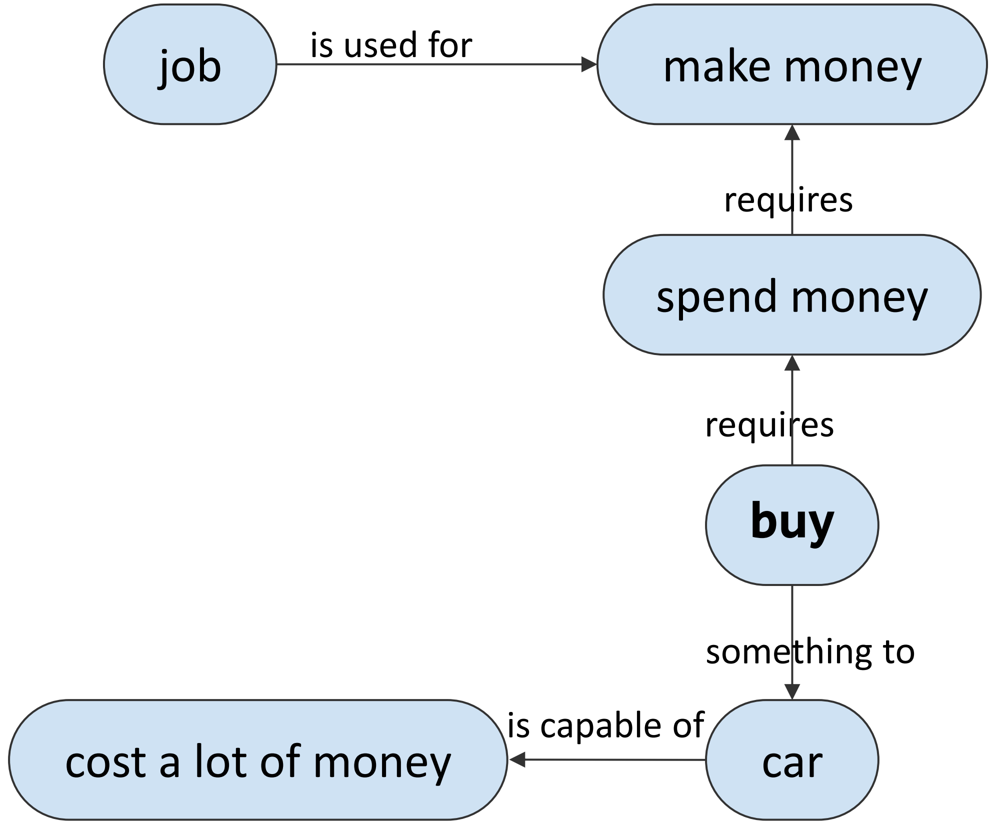

Knowledge Source

Conceptnet 5.5: An open multilingual graph of general knowledge. *Robyn Speer, Joshua Chin, and Catherine Havasi.* AAAI 2017.

Other Knowledge Sources

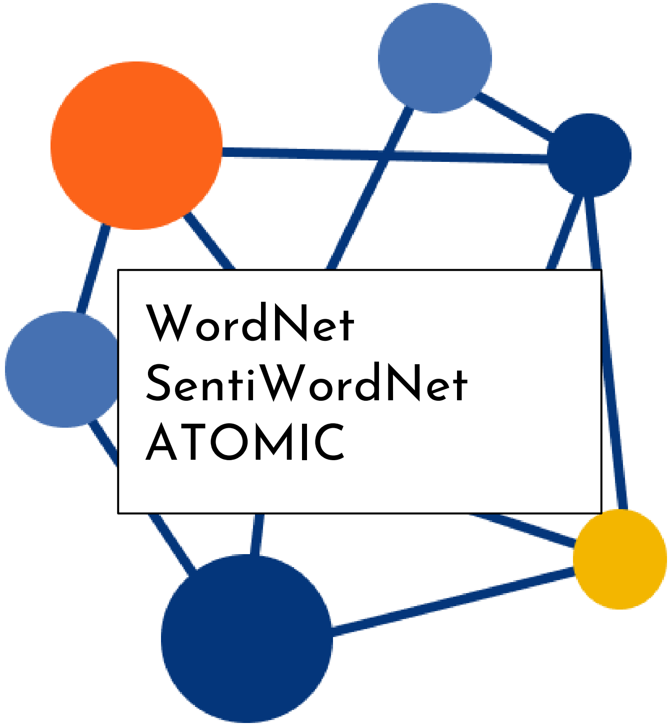

WordNet
SentiWordNet
ATOMIC

Mining from Text

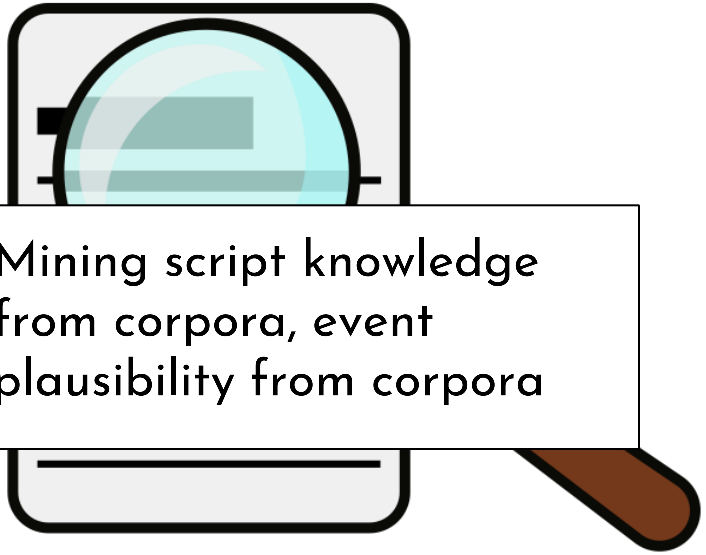

Mining script knowledge from corpora, event plausibility from corpora

Knowledge Bases

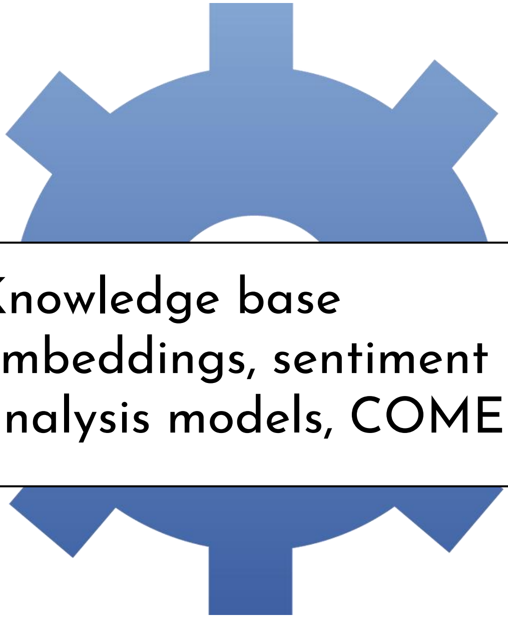

Knowledge base embeddings, sentiment analysis models, COMET

Tools

Neural Component

[CLS] Katrina had the financial means to afford a new car while Monica did not, since [SEP] Katrina had a high paying job.

[CLS] Katrina had the financial means to afford a new car while Monica did not, since [SEP] Monica had a high paying job.

0.51
0.49

Neural

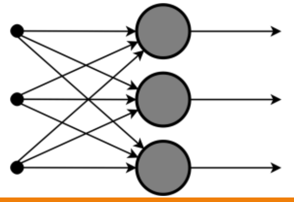

COMET's Combination Method

Incorporate into scoring function

Symbolic → vector representation

Multi-task learning

- (This was before we had *very large* LMs)

Combined

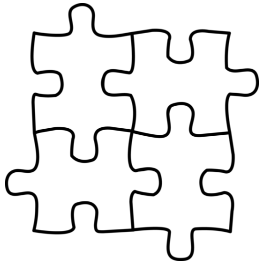

Incorporating External Knowledge into Neural Models

Example

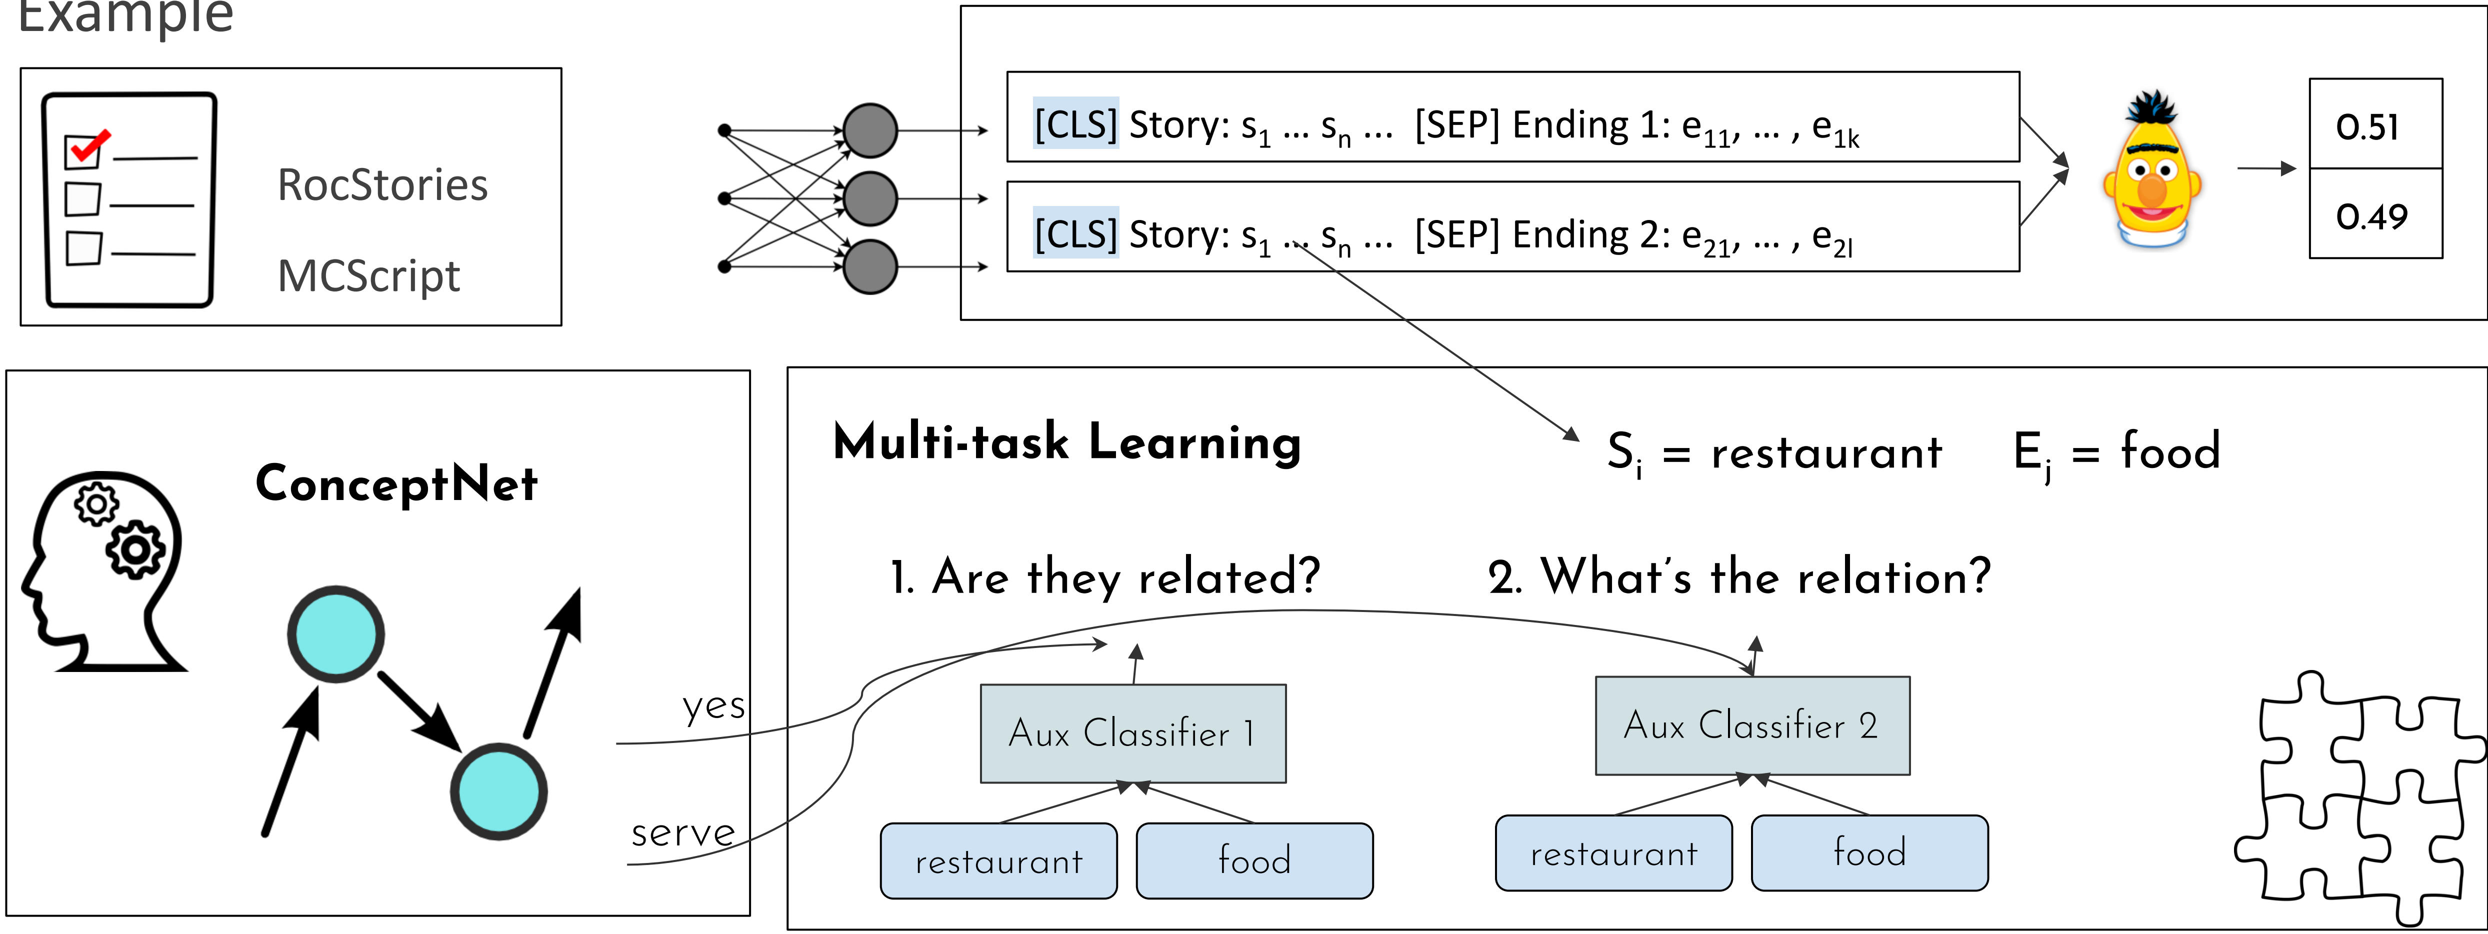

Incorporating Commonsense Reading Comprehension with Multi-task Learning. *Jiangnan Xia, Chen Wu, and Ming Yan.* CIKM 2019.

Review: Limitations

- Insufficient Coverage
- Not 100% accurate

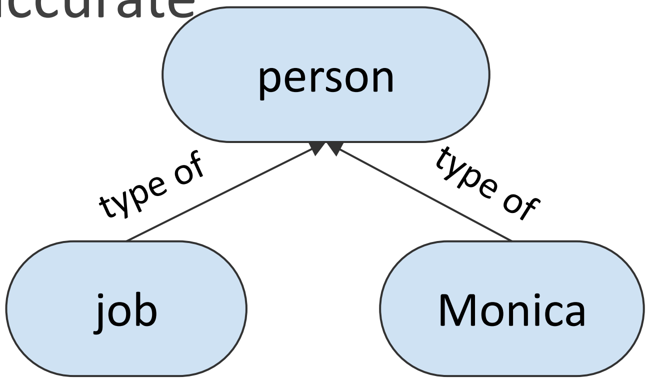

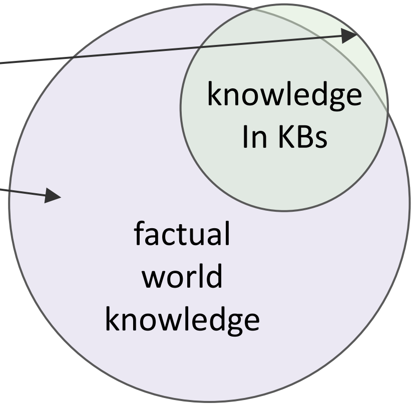

- Easy to incorporate simple resources with stationary facts (ConceptNet) but they are limited in expressiveness:

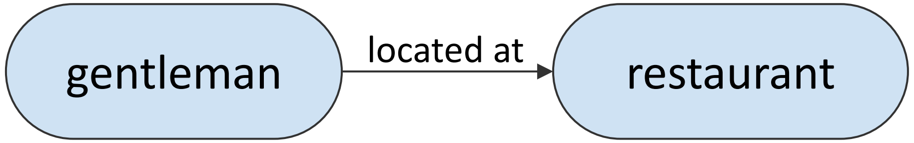

Limitations of Knowledge Graphs

- Situations rarely found as-is in commonsense knowledge graphs

ATOMIC

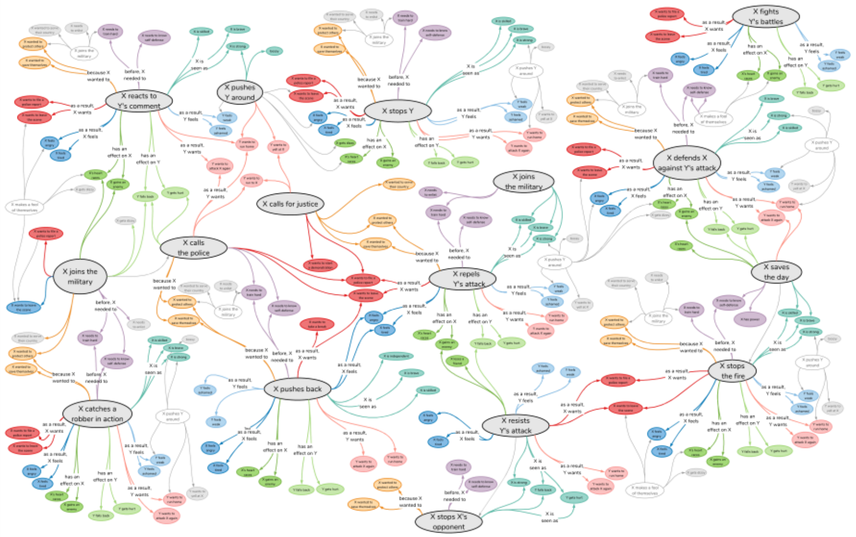

(Sap et al., 2019)

(X goes to the mall, Effect on X, buys clothes)

(X goes the mall, Perception of X, rich)

(X gives Y some money, Reaction of Y, grateful)

Limitations of Knowledge Graphs

Kai knew that things were getting out of control and managed to keep his temper in check

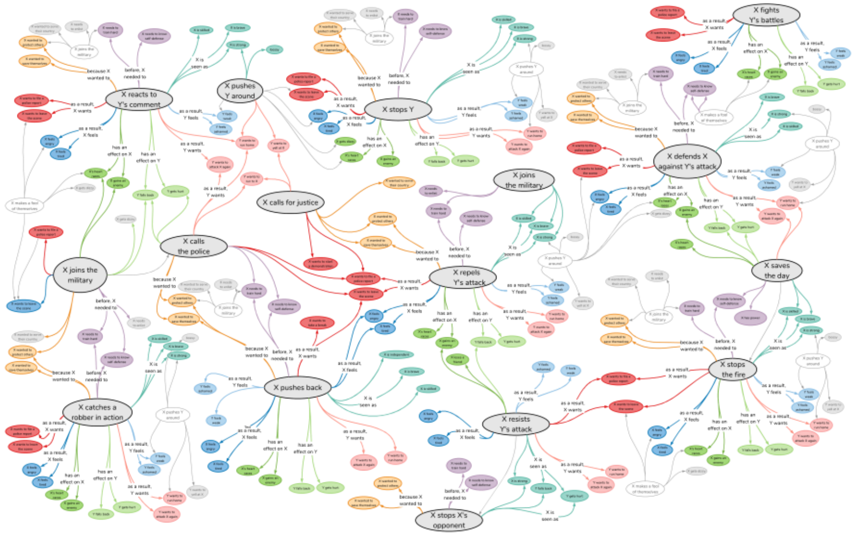

(Sap et al., 2019)

Limitations of Knowledge Graphs

- Situations rarely found as-is in commonsense knowledge graphs
- Connecting to knowledge graphs can yield incorrect nodes

Kai knew that things were getting out of control and managed to keep his temper in check

X keeps X's temper

X keeps ___ under control

X keeps X's ___ in check

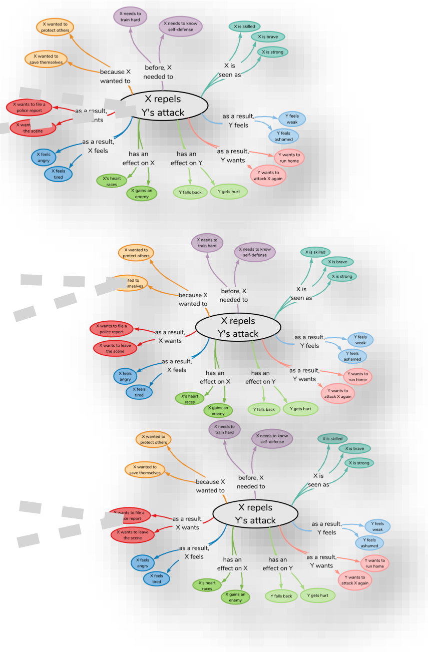

Limitations of Knowledge Graphs

- Situations rarely found as-is in commonsense knowledge graphs
- Connecting to knowledge graphs can yield incorrect nodes
- Suitable nodes are often uncontextualized

Kai knew that things were getting out of control and managed to keep his temper in check

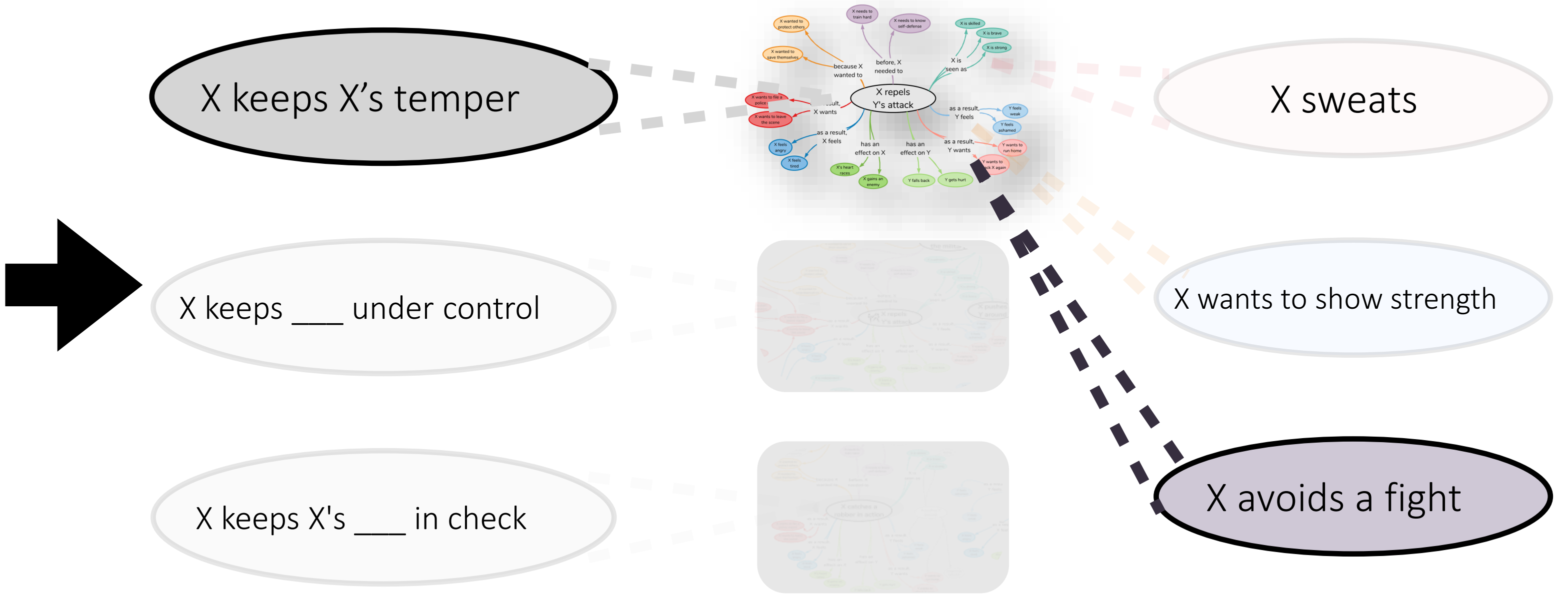

Challenge

How do we provide machines with large-scale commonsense knowledge?

Constructing Knowledge Graphs

Observe world

Write commonsense knowledge facts

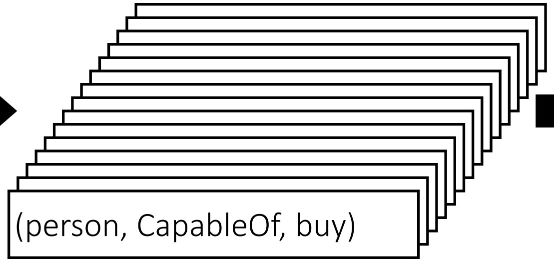

Store facts in knowledge graph

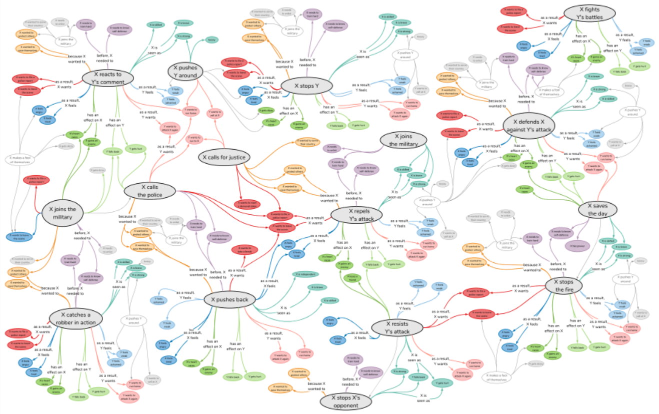

Constructing Knowledge Graphs

Observe world

Write commonsense knowledge facts

Store facts in knowledge graph

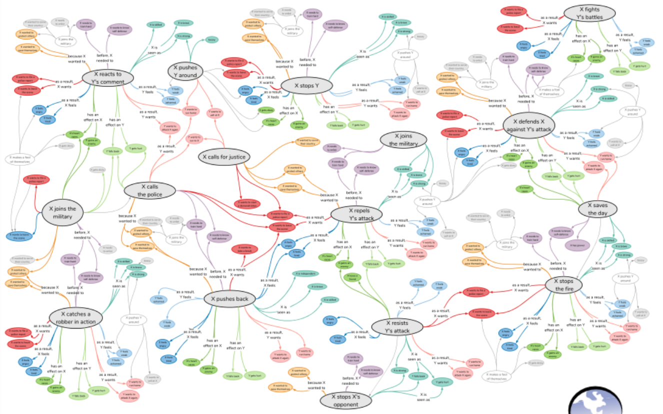

(Miller, 1995)

(Singh et al., 2002)

(Lenat, 1995)

(Sap et al., 2019)

Challenges of Prior Approaches

Commonsense knowledge is immeasurably vast, making it impossible to manually enumerate

Constructing Knowledge Graphs Automatically

Gather Textual Corpus

John went to the grocery **store** to buy some **steaks**. He was going to prepare **dinner** for his **daughter**'s **birthday**. She was turning 5 and would be starting elementary **school** soon.

Automatically extract knowledge

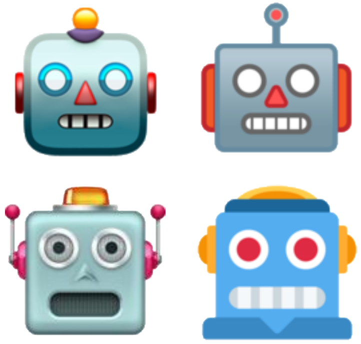

(Schubert, 2002)

(Banko et al., 2007)

(Zhang et al., 2020)

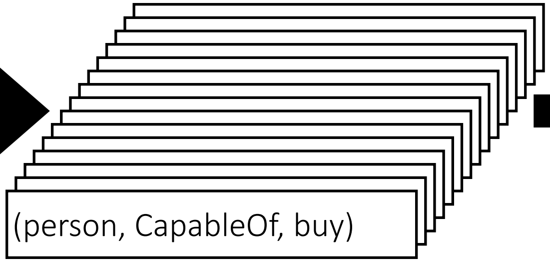

Store in knowledge graph

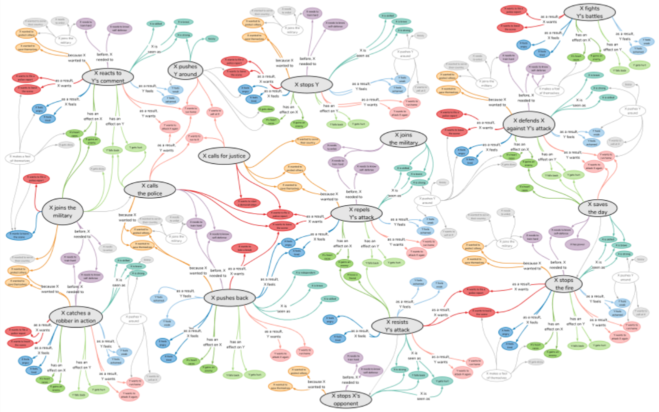

ConceptNet
An open, multilingual knowledge graph

(Speer et al., 2017)

Webchild

(Tandon et al., 2019)

Challenges of Prior Approaches

Commonsense knowledge is immeasurably vast, making it impossible to manually enumerate

Commonsense knowledge is often implicit, and often can’t be directly extracted from text

Knowledge Base Completion

Gather training set of knowledge tuples

(person, CapableOf, buy)

Learn relationships among entities

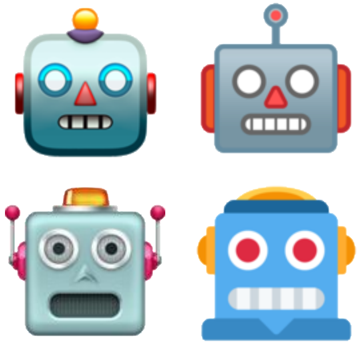

(Socher et al., 2013)

(Bordes et al., 2013)

(Riedel et al., 2013)

(Toutanova et al., 2015)

(Yang et al., 2015)

(Trouillon et al., 2016)

(Nguyen et al., 2016)

(Dettmers et al., 2018)

Predict new relationships

(person, CapableOf,)?

Store in knowledge graph

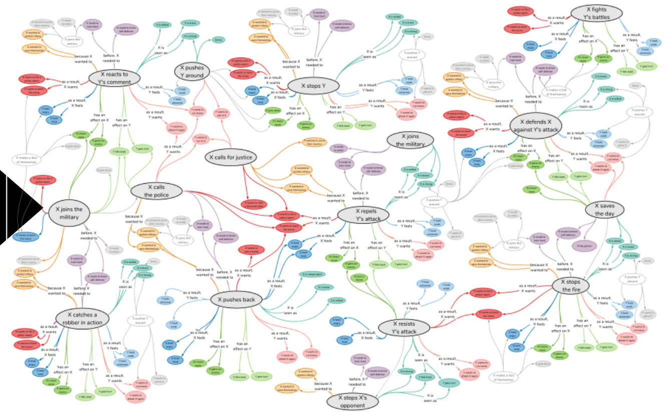

Commonsense Knowledge Base Completion

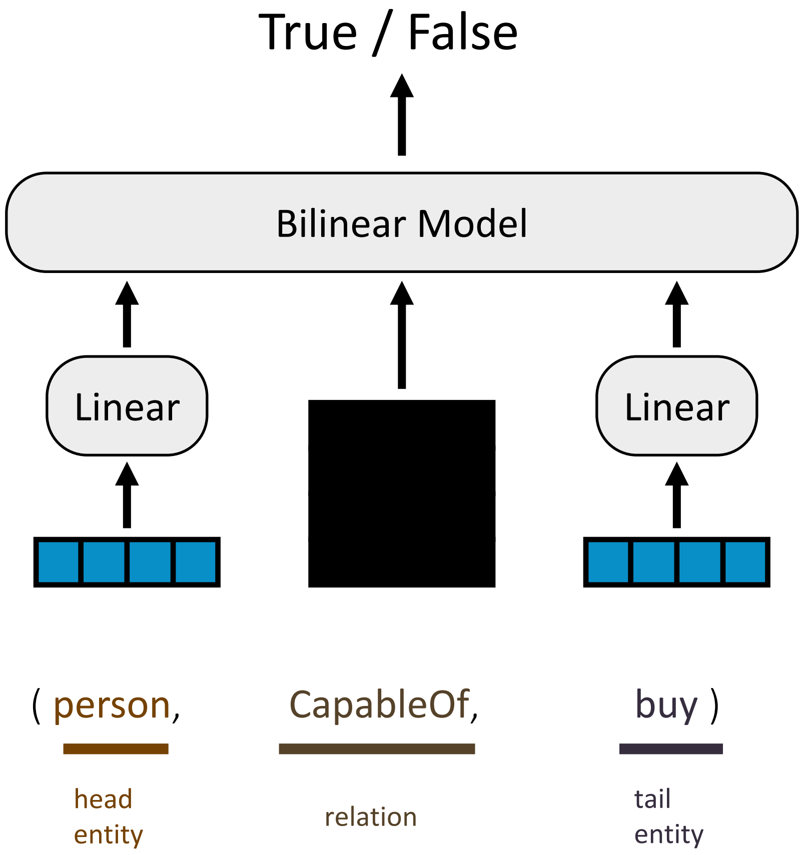

Li et al., 2016

Only high confidence predictions are validated

Low Novelty

Jastrzebski et al., 2018

Commonsense Knowledge Base Completion and Generation!

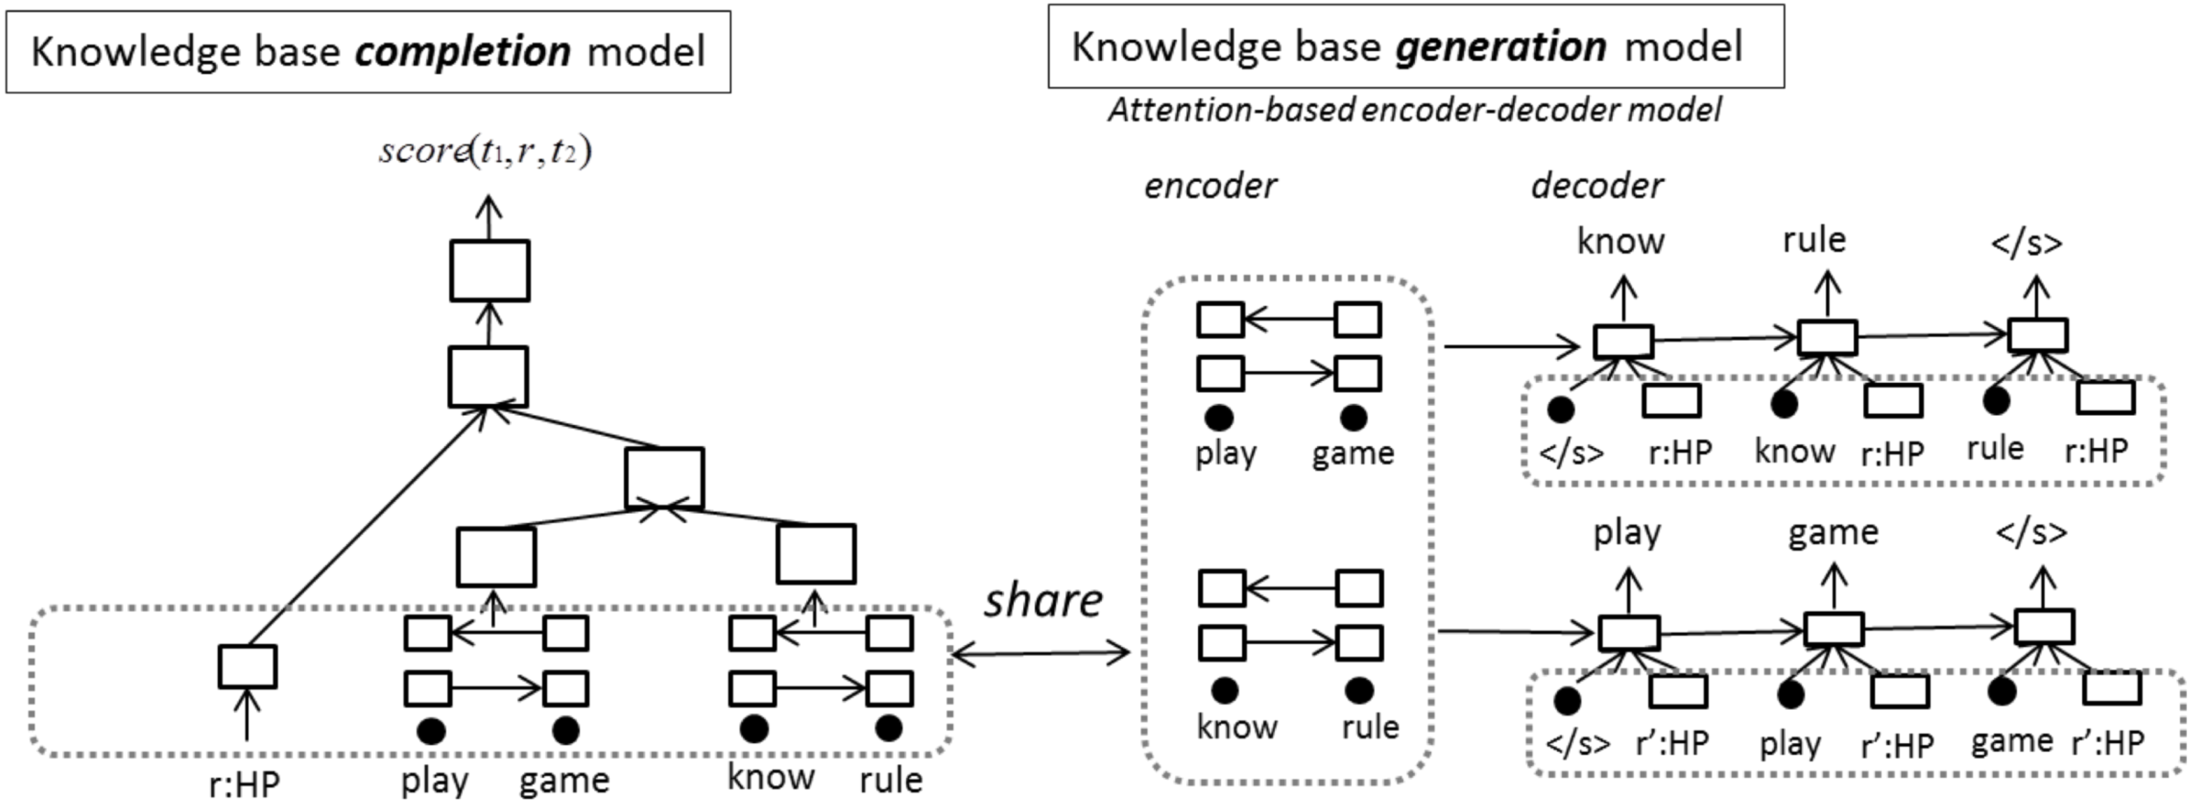

Challenges of Prior Approaches

Commonsense knowledge is immeasurably vast, making it impossible to manually enumerate

Commonsense knowledge is often implicit, and often can’t be directly extracted from text

Commonsense knowledge resources are quite sparse, making them difficult to extend by only learning from examples

Solution Outline

Leverage manually curated commonsense knowledge resources

Learn from the examples to induce new relationships

Scale up using language resources

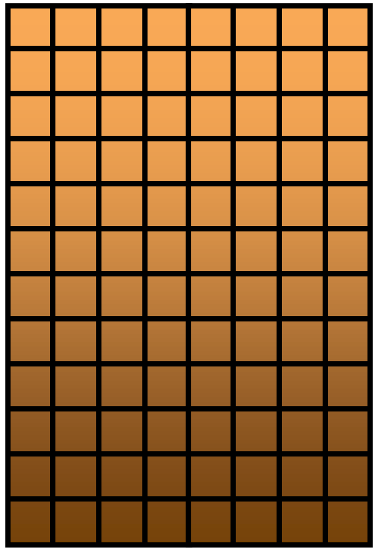

Learn word embeddings from language corpus

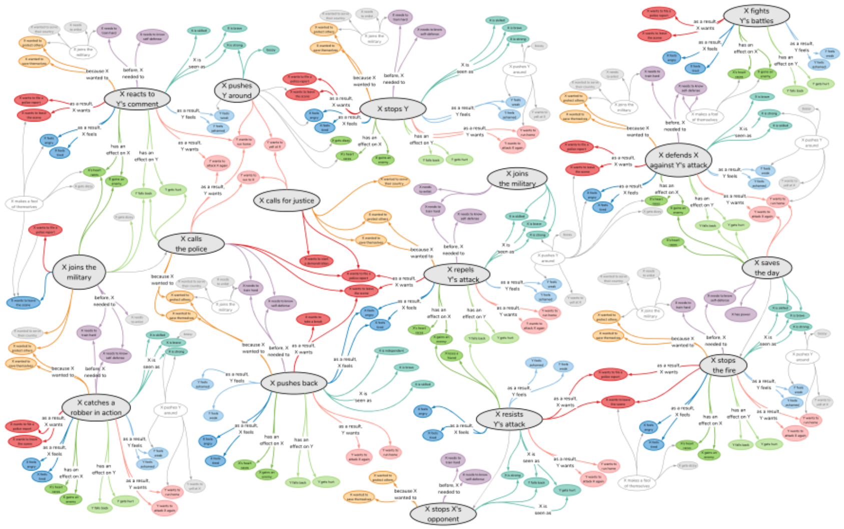

Retrofit word embeddings on semantic resource

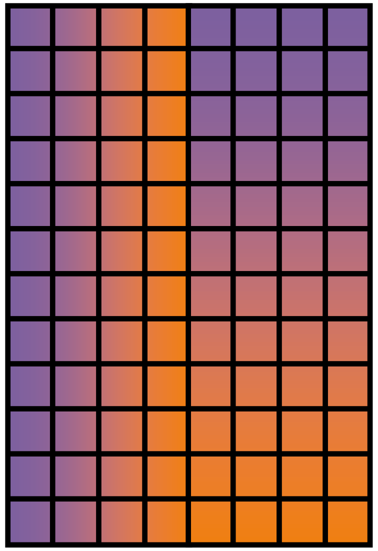

Learn knowledge-aware embeddings

Faruqui et al., 2015, Speer et al., 2017

Structure of Knowledge Tuple

Learning Structure of Knowledge

Given a seed entity and a relation, learn to generate the target entity

$$\mathcal{L} = -\sum \log P(\textit{target words} \mid \textit{seed words, relation})$$

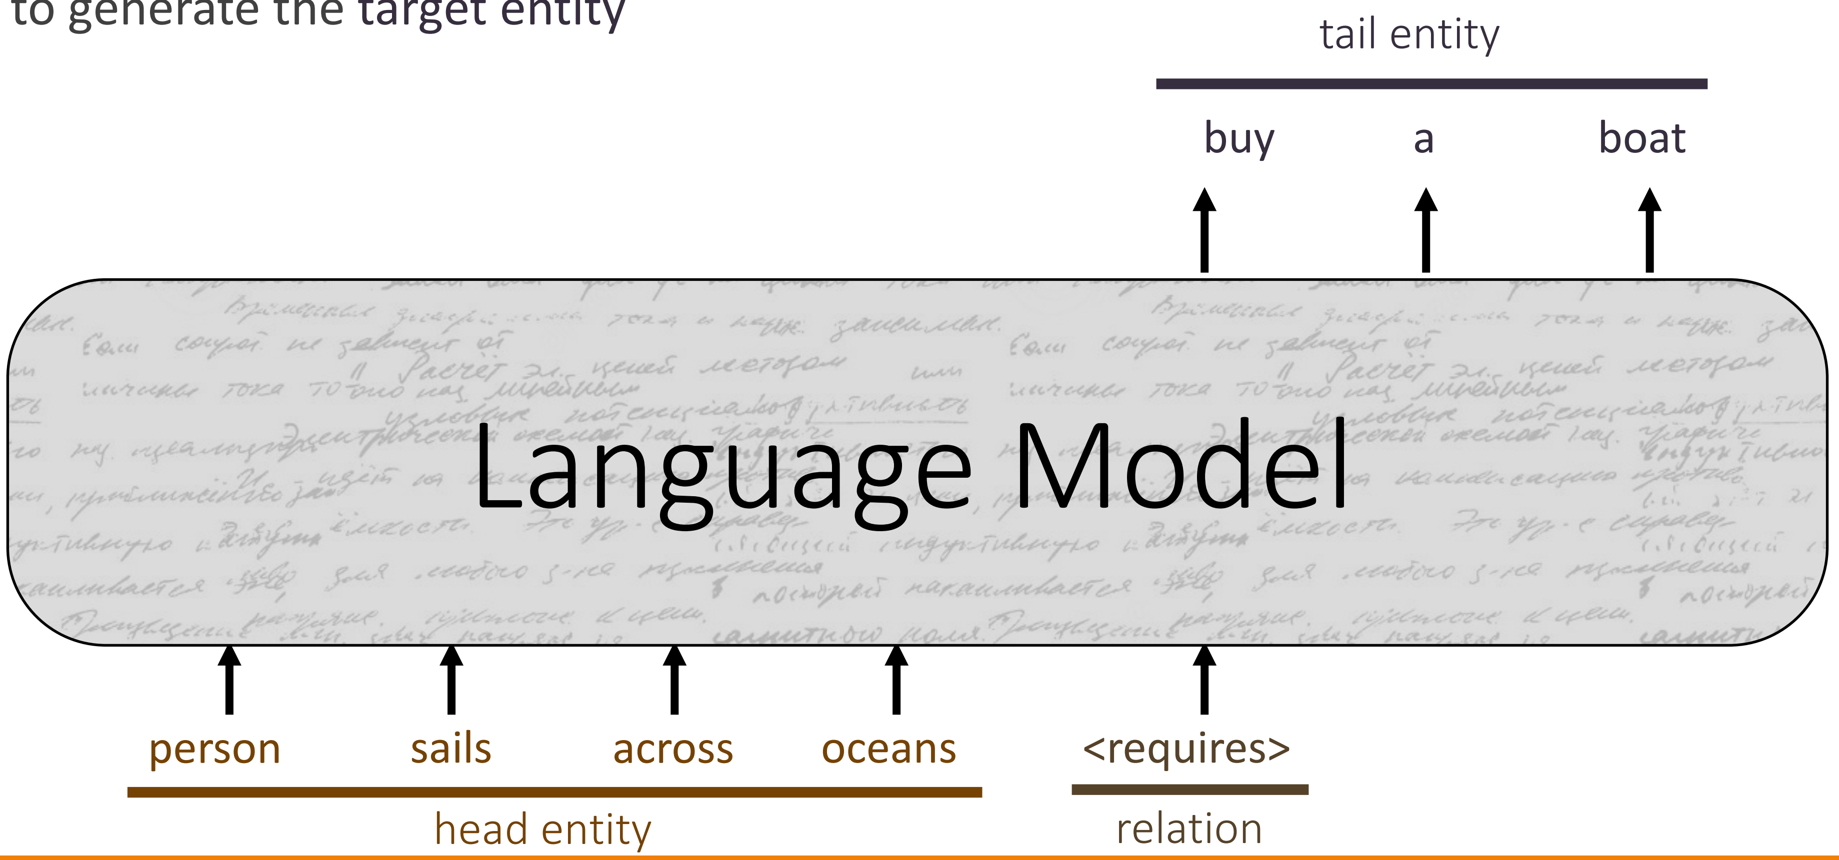

(Bosselut et al., 2019)

Learning Structure of Knowledge

Language Model → Knowledge Model: generates knowledge of the structure of the examples used for training

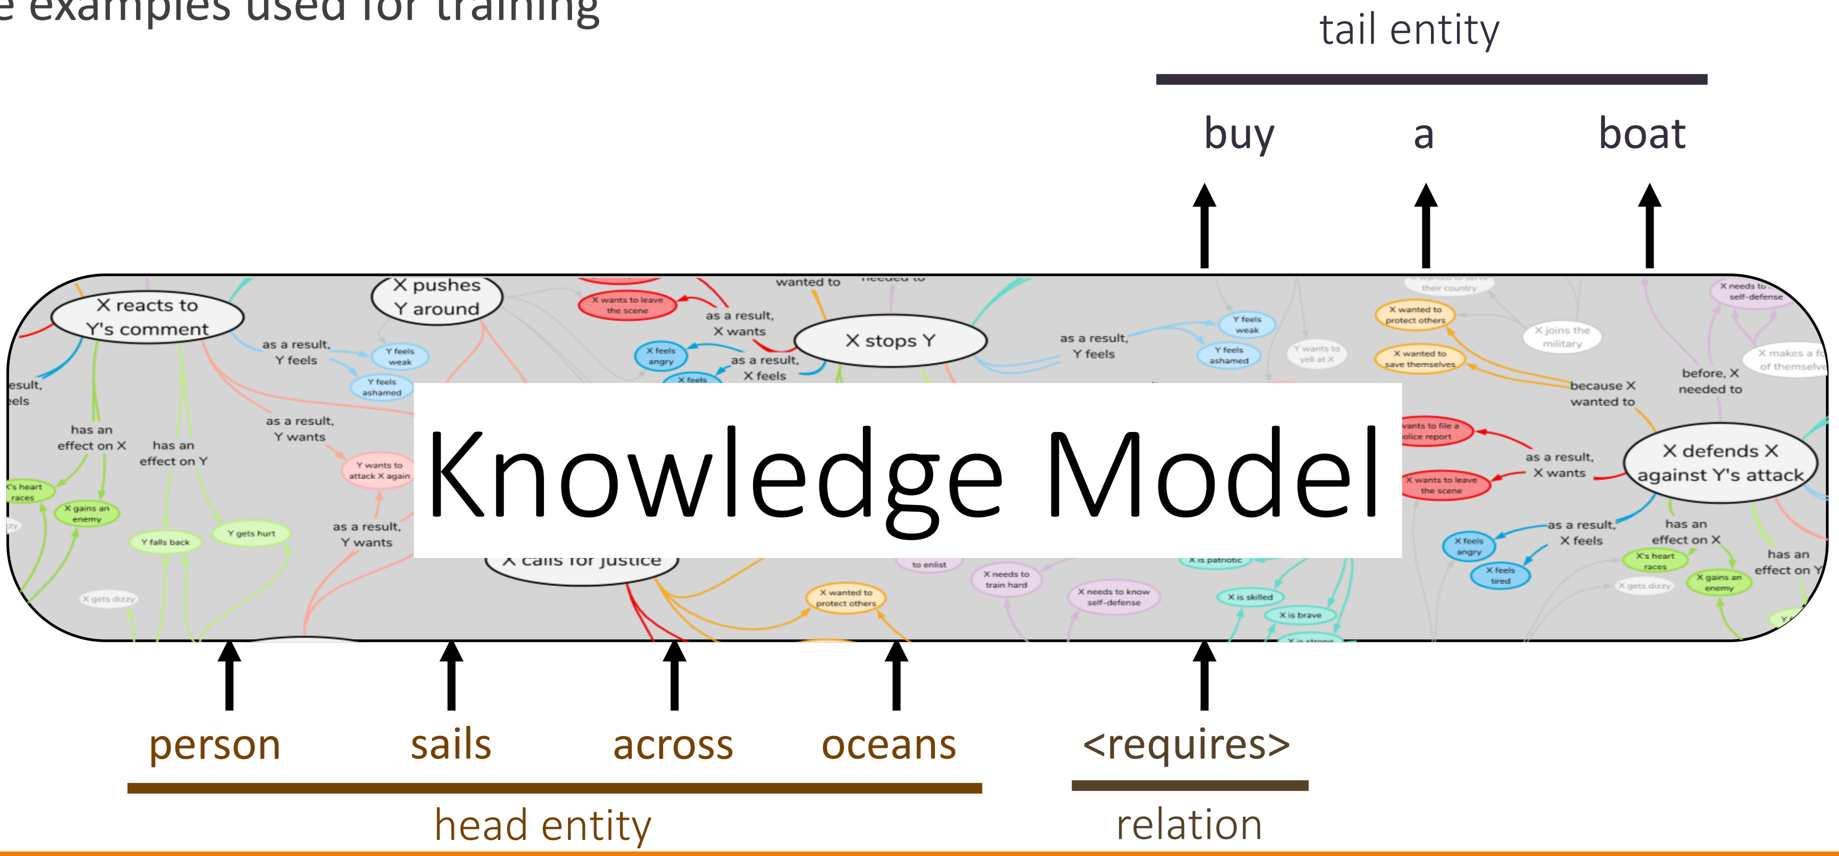

(Bosselut et al., 2019)

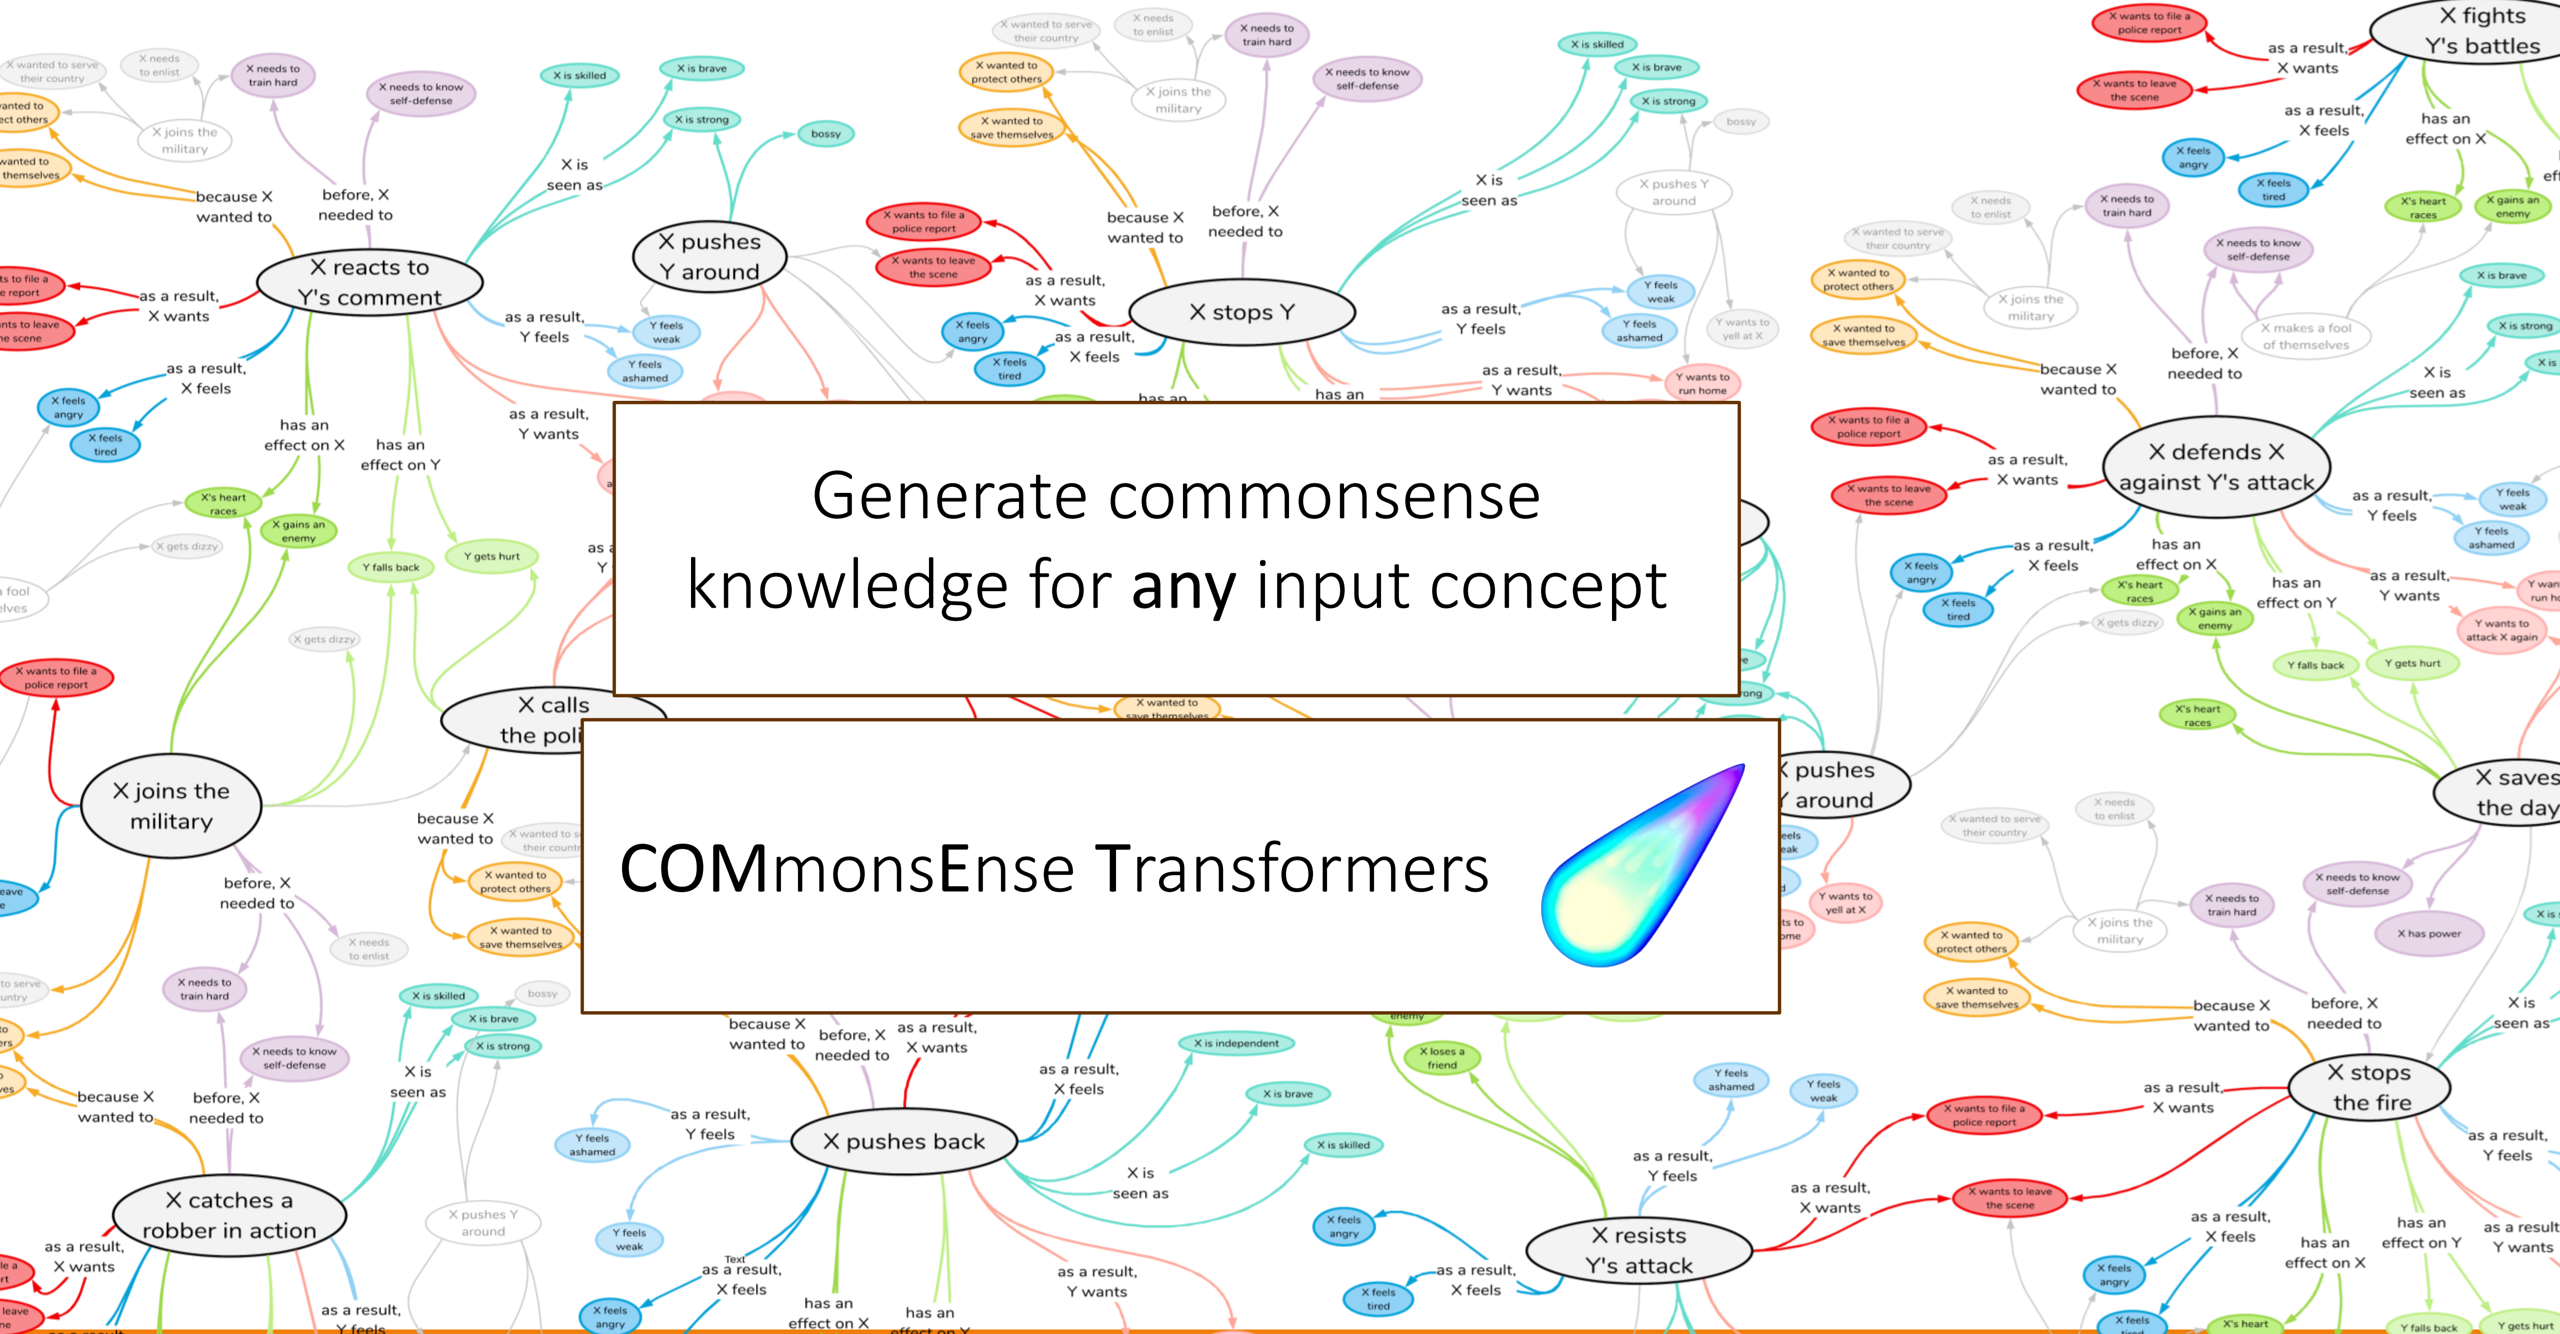

Generate commonsense knowledge for **any** input concept

COMmons**E**nse **T**ransformers

COMET - ATOMIC

PersonX gives a tutorial

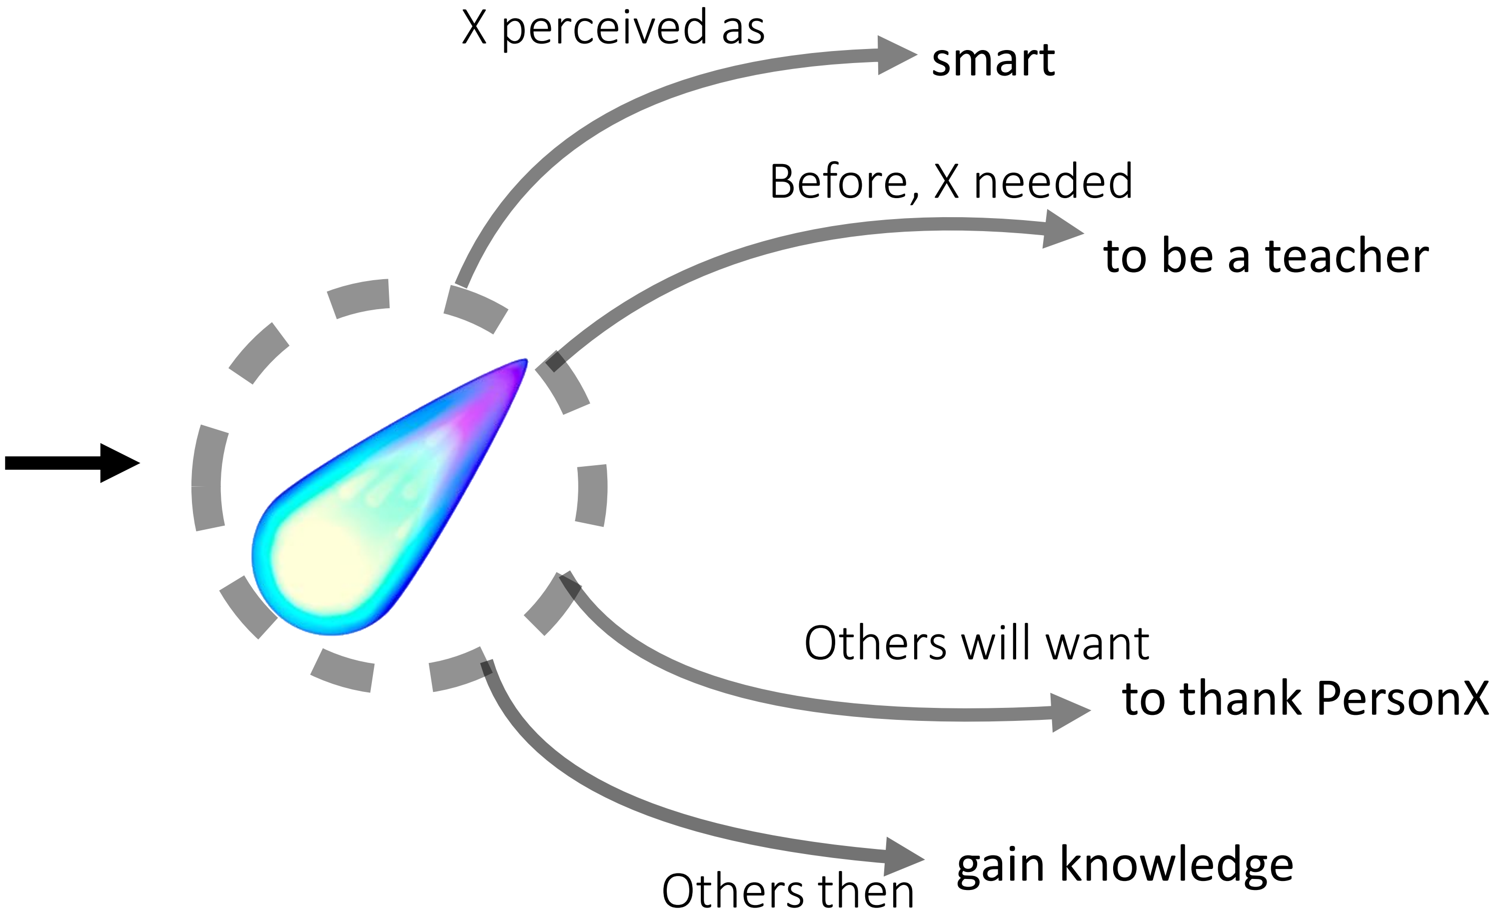

COMET - ConceptNet

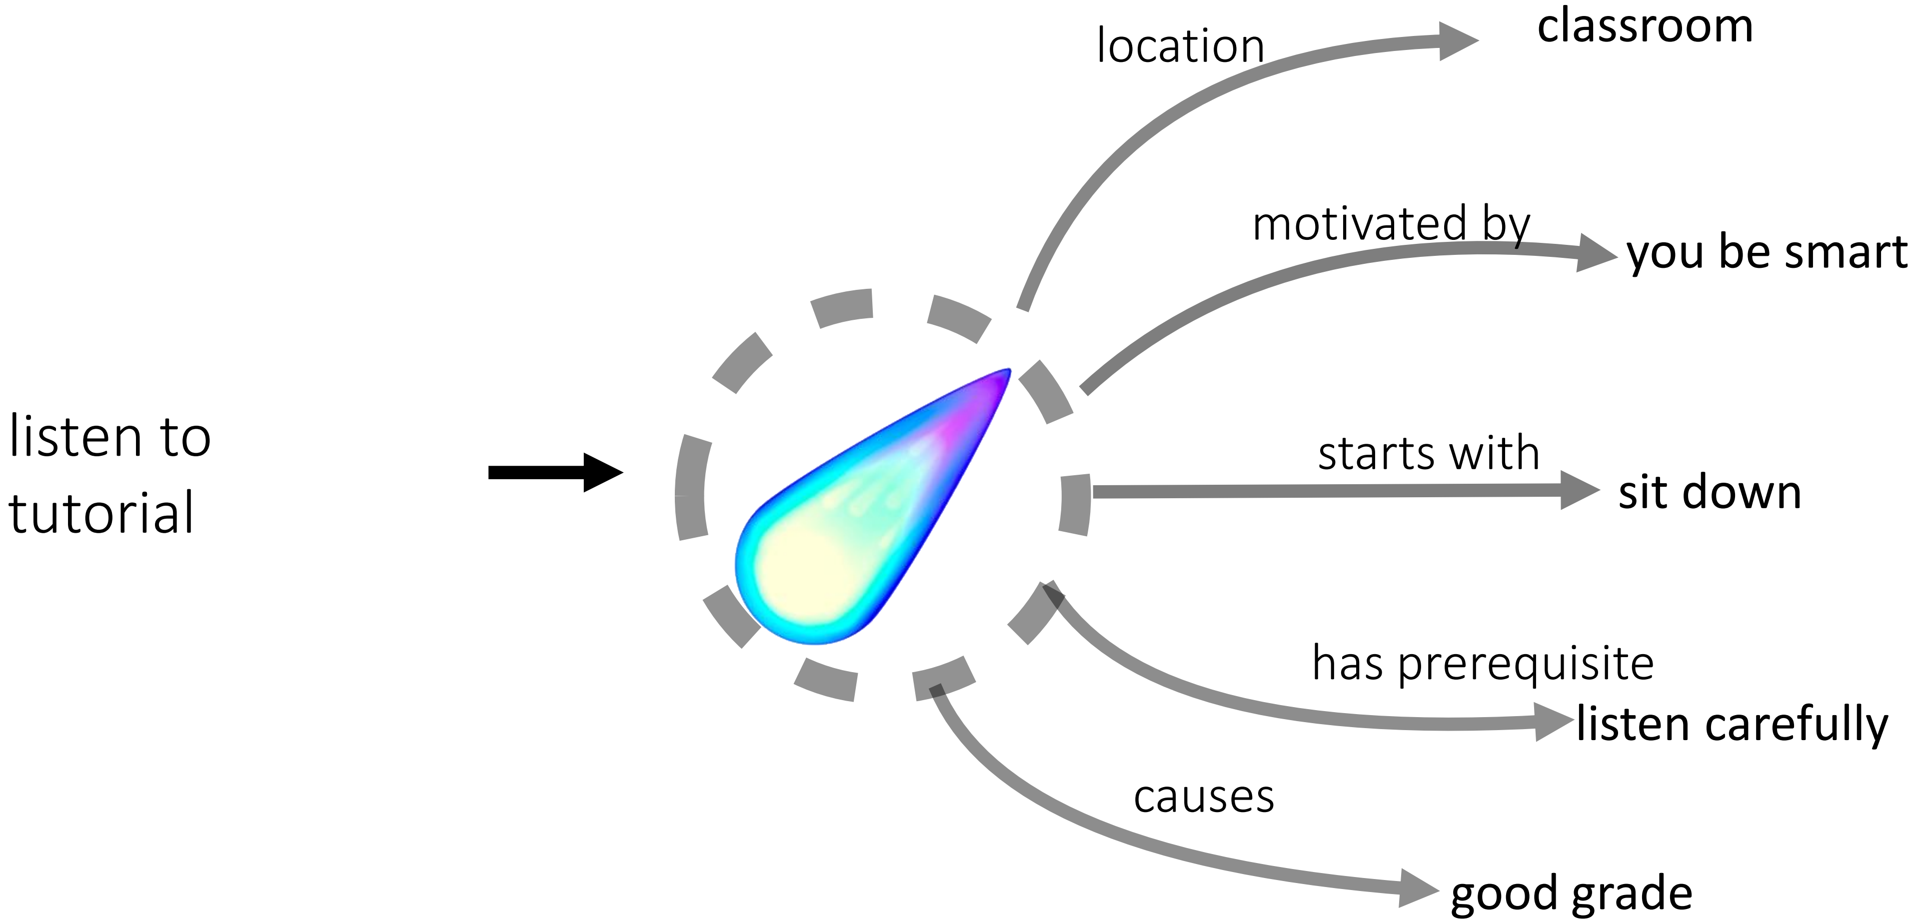

Why does this work?

Transfer Learning from Language

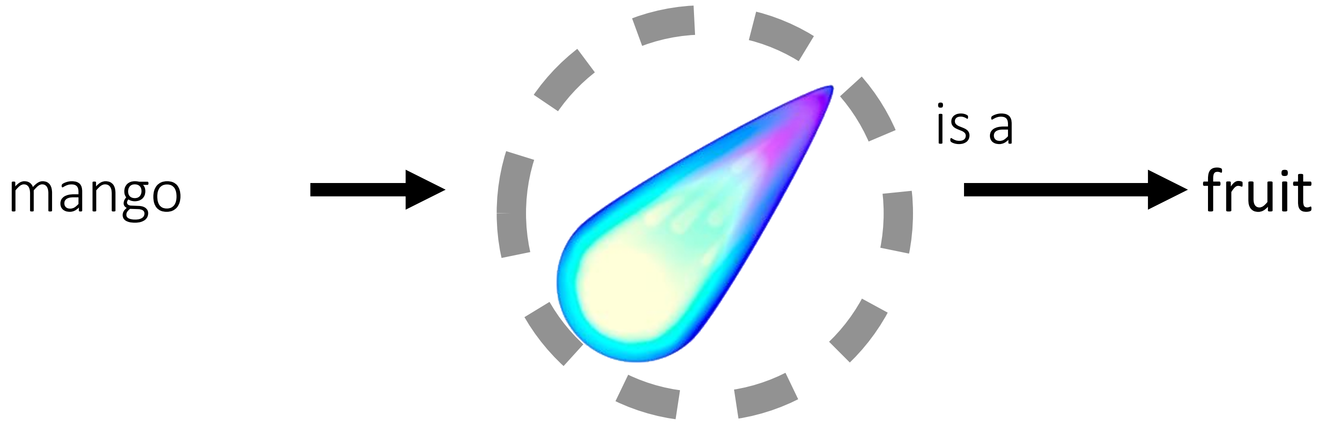

Transfer Learning from Language

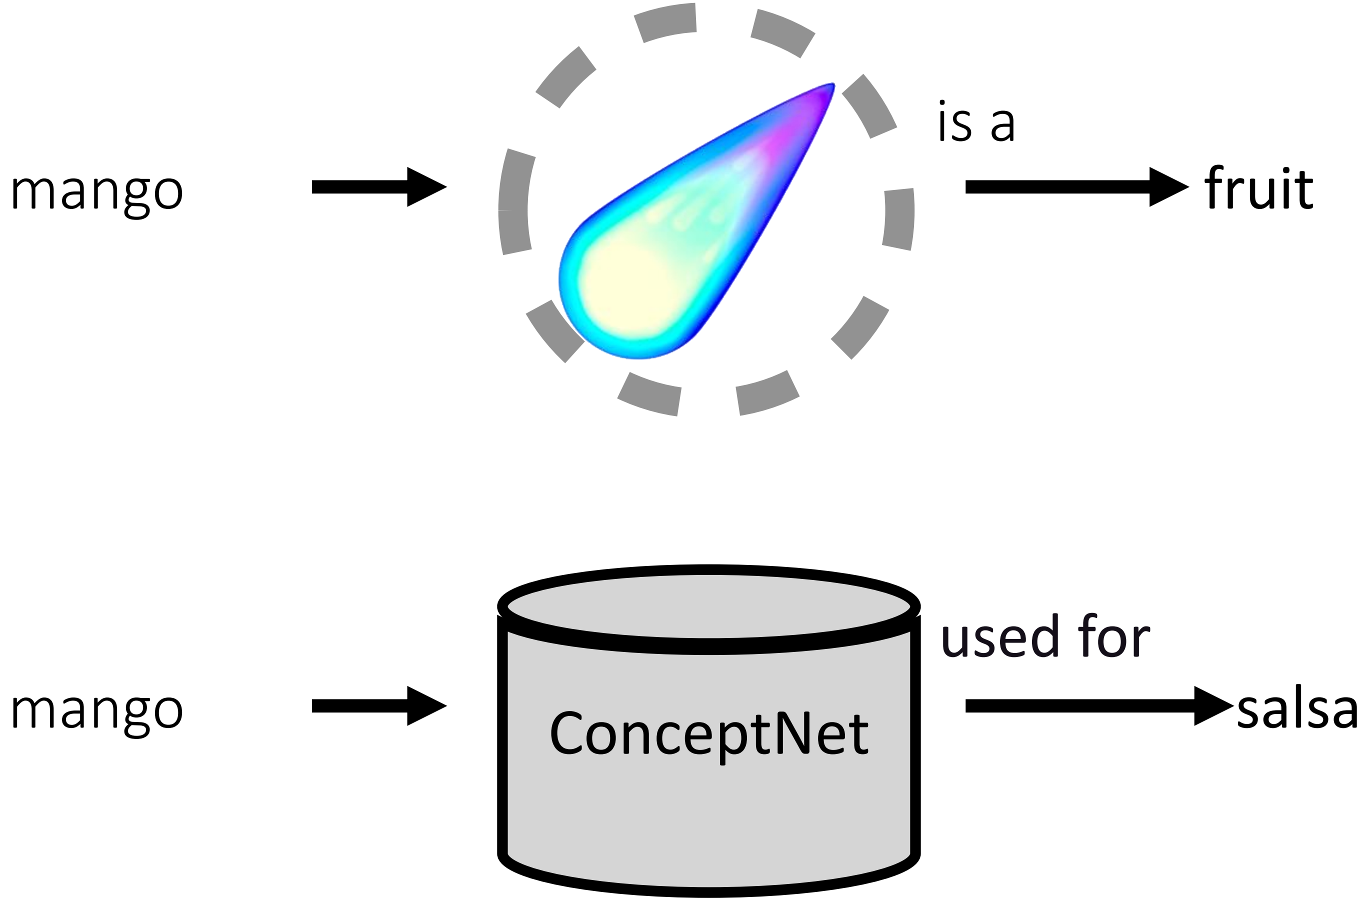

Transfer Learning from Language

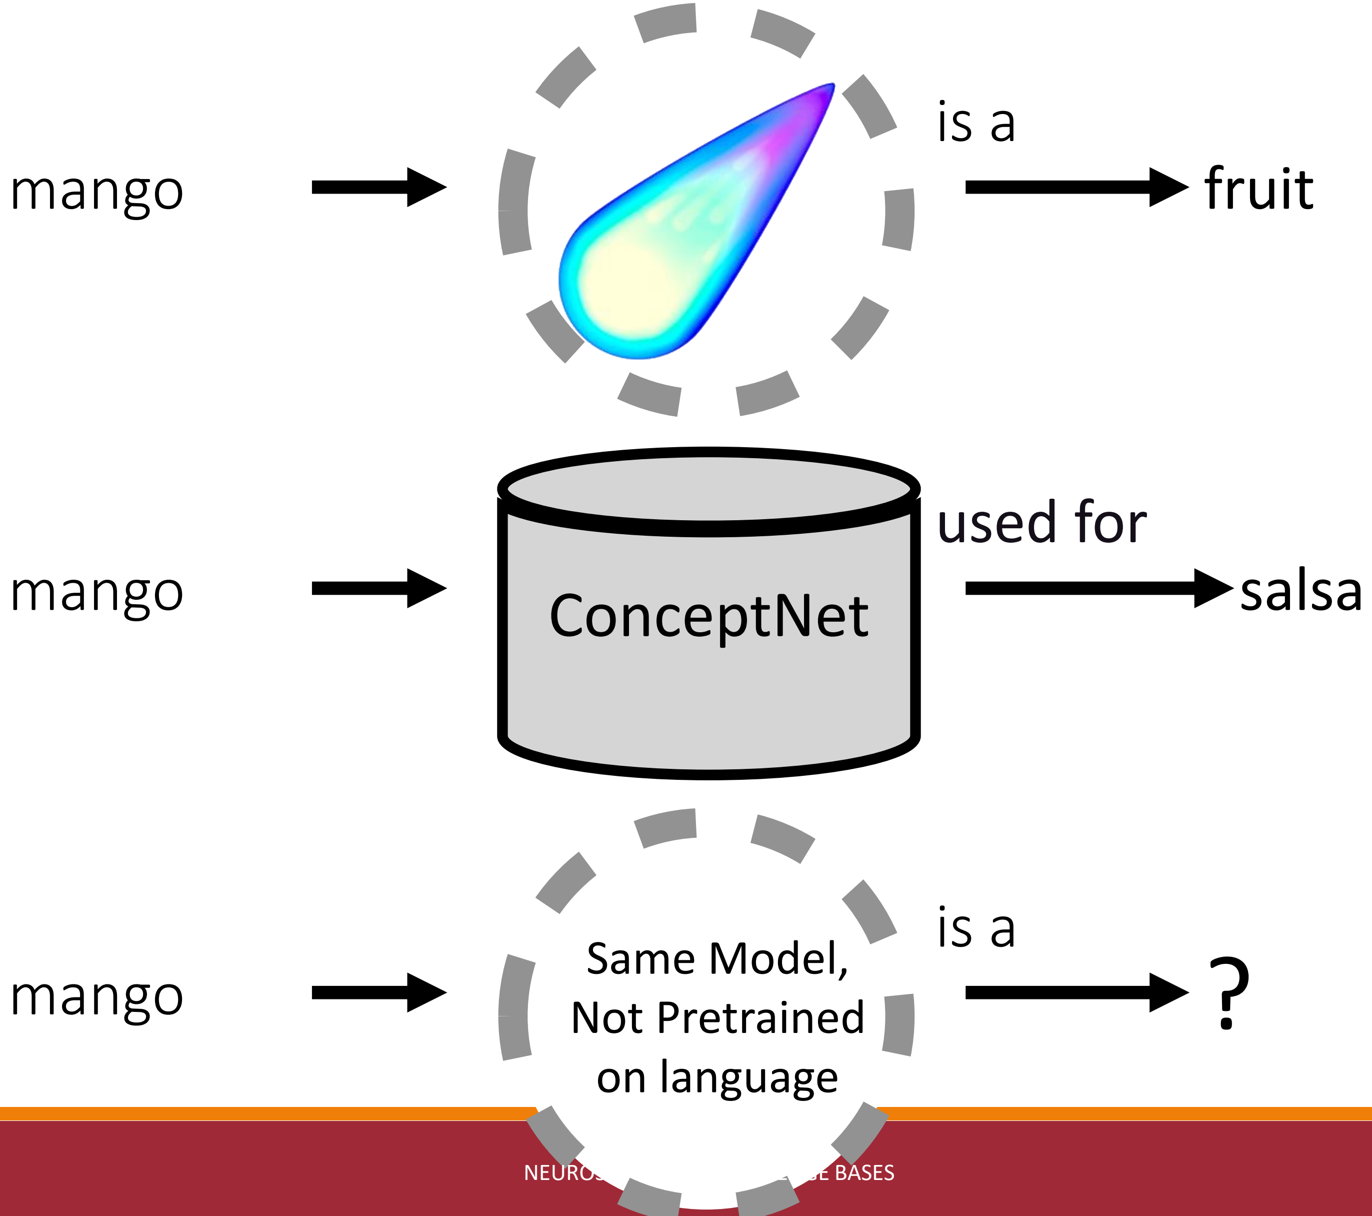

Transfer Learning from Language

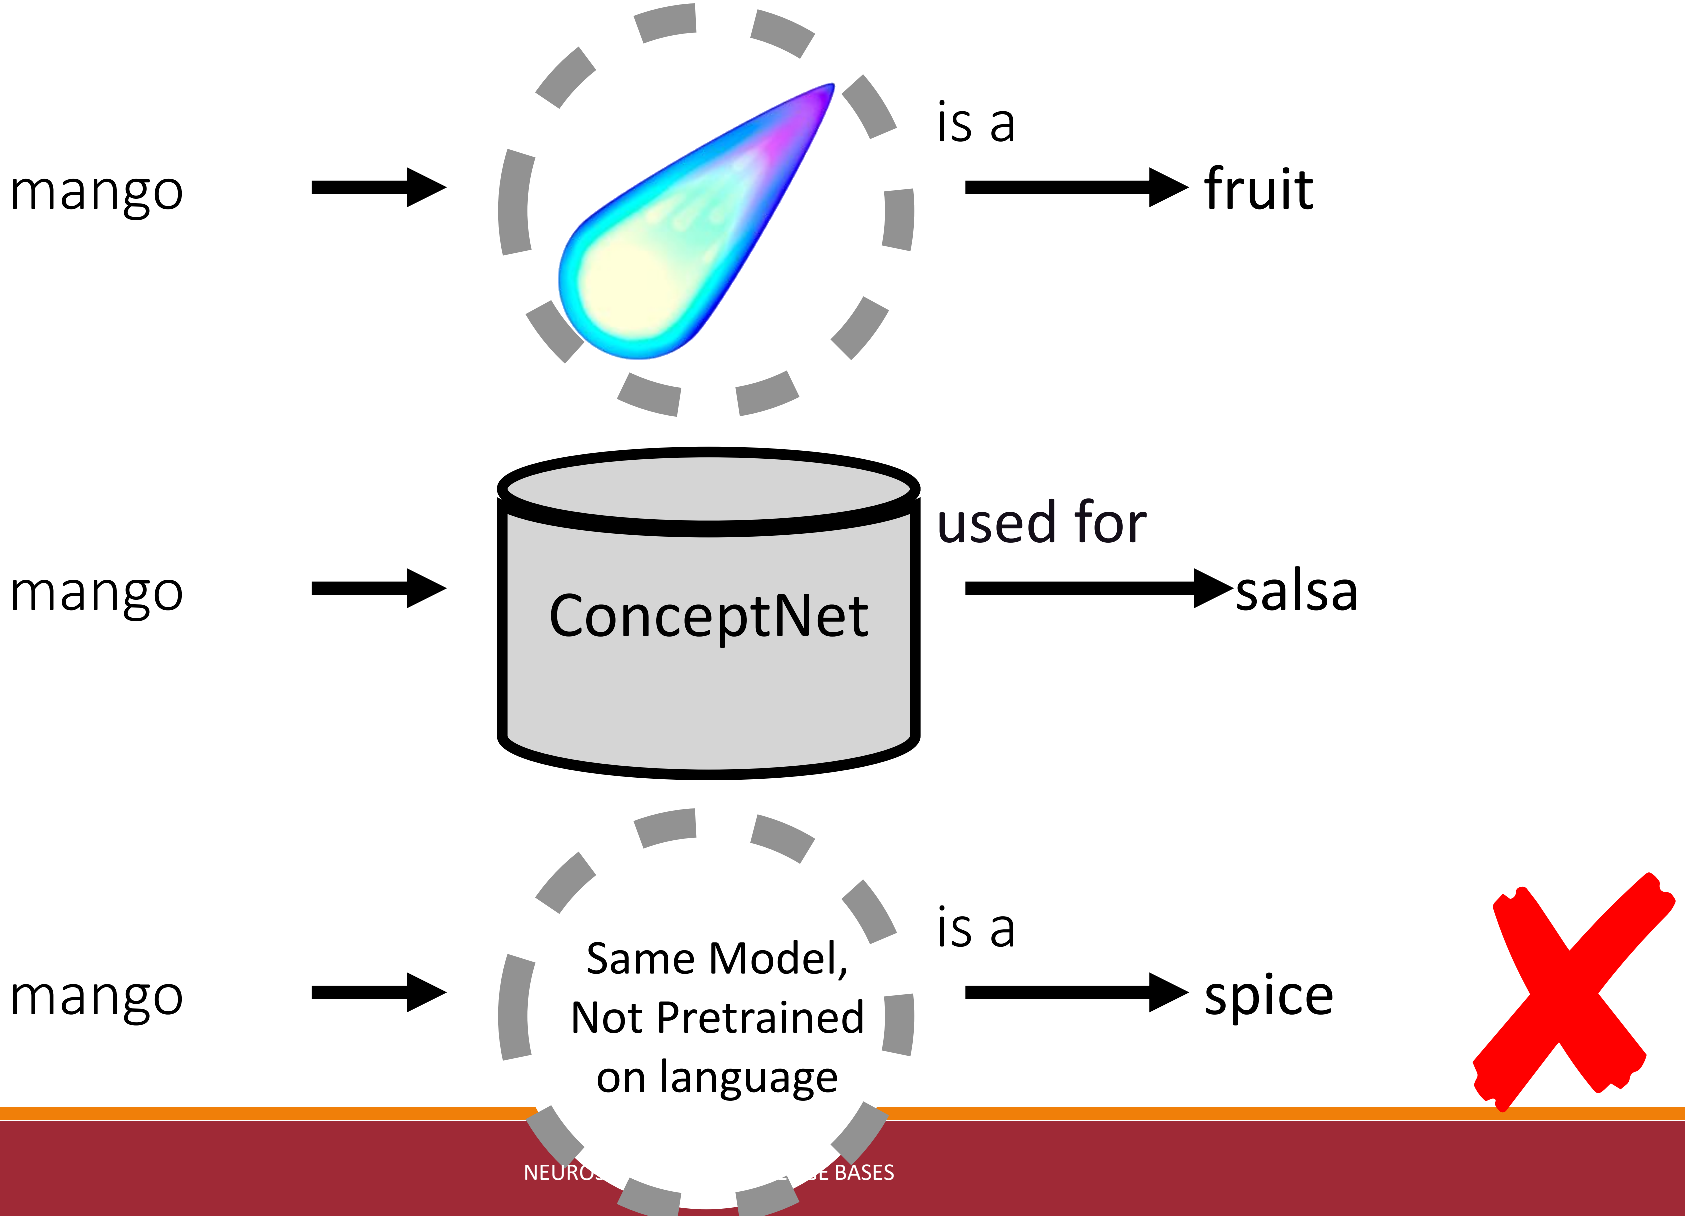

Can’t an off-the-shelf language model do the same thing?

Do Language Models know this?

Sentence:

mango is a

Predictions:

2.1% **great**
1.9% **very**
1.2% **new**
1.0% **good**
1.0% **small**
← **Undo**

https://demo.allennlp.org/next-token-lm

Do Language Models know this?

Sentence:

mango is a

Predictions:

- 2.1% great
- 1.9% very
- 1.2% new
- 1.0% good
- 1.0% small
- ← Undo

Sentence:

a mango is a

Predictions:

- 4.2% good
- 4.0% very
- 2.5% great
- 2.4% delicious
- 1.8% sweet
- ← Undo

Do Language Models know this?

Sentence:

mango is a

Predictions:

2.1% great
1.9% very
1.2% new
1.0% good
1.0% small
← Undo

Sentence:

a mango is a

Predictions:

4.2% good
4.0% very
2.5% great
2.4% delicious
1.8% sweet
← Undo

Sentence:

A mango is a

Predictions:

4.2% fruit
3.5% very
2.5% sweet
2.2% good
1.5% delicious
← Undo

Do *Masked* Language Models know this?

Sentence:

mango is a [MASK]

Mask 1 Predictions:

- 69.7% **.**
- 9.3% **;**
- 1.7% **!**
- 0.8% **vegetable**
- 0.7% **?**

Sentence:

mango is a [MASK] .

Mask 1 Predictions:

- 7.6% **staple**
- 7.6% **vegetable**
- 4.6% **plant**
- 3.5% **tree**
- 3.5% **fruit**

Sentence:

A mango is a [MASK] .

Mask 1 Predictions:

- 16.0% **banana**
- 12.1% **fruit**
- 5.9% **plant**
- 5.5% **vegetable**
- 2.5% **candy**

Think-Pair-Share

How would you get a modern LM to produce the correct behavior without finetuning?

Sensitivity to cues

Candidate Sentence S_i	$\log p(S_i)$
"musician can playing musical instrument"	-5.7
"musician can be play musical instrument"	-4.9
"musician often play musical instrument"	-5.5
"a musician can play a musical instrument"	$-\mathbf{2.9}$

Feldman et al., 2019

Prompt	**Model Predictions**
A ___ has fur.	dog, cat, fox, ...
A ___ has fur, is big, and has claws.	cat, **bear**, lion, ...
A ___ has fur, is big, has claws, has teeth, is an animal, eats, is brown, and lives in woods.	**bear**, wolf, cat, ...

Weir et al., 2020

Commonsense Transformers

- Language models **implicitly represent knowledge**

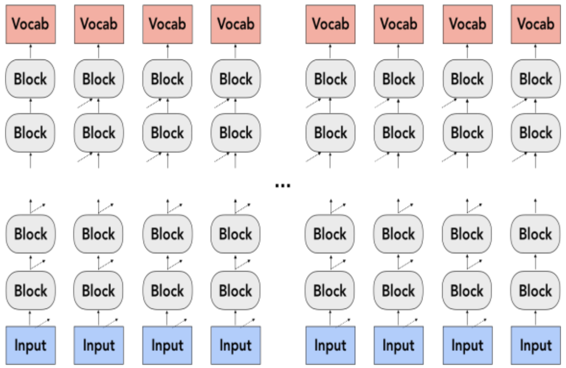

Pre-trained
Language Model

Commonsense Transformers

- Language models **implicitly represent knowledge**
- Finetune them on knowledge graphs to **learn structure of knowledge**

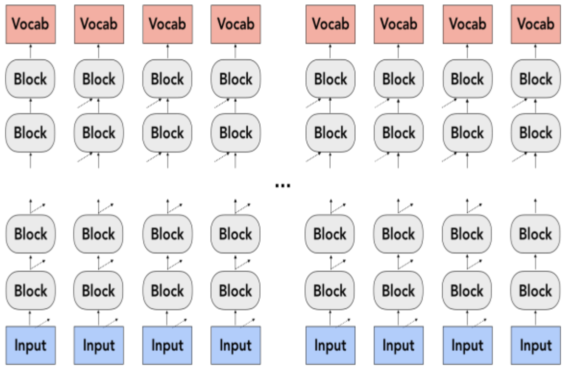

+

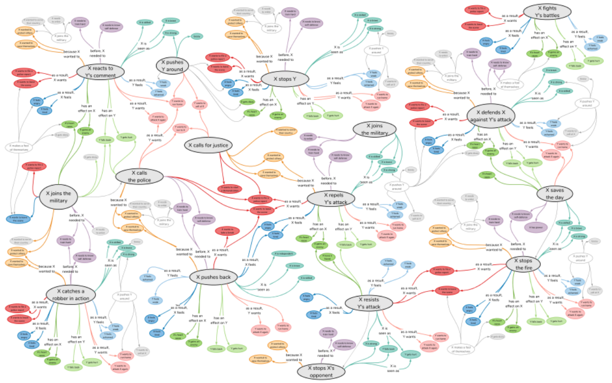

Pre-trained Language Model

Seed Knowledge Graph Training

Commonsense Transformers

- Language models **implicitly represent knowledge**
- Finetune them on knowledge graphs to **learn structure of knowledge**
- Resulting knowledge model **generalizes structure** to other concepts

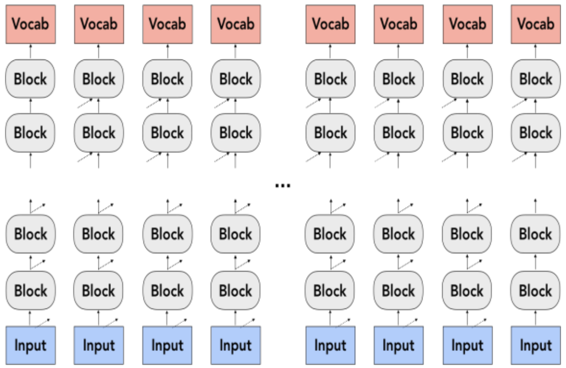

Pre-trained
Language Model

+

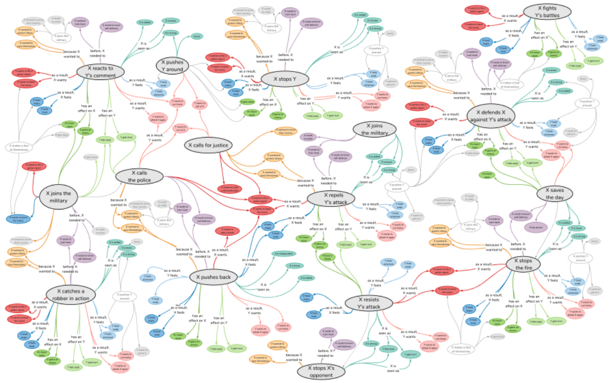

Seed Knowledge
Graph Training

=

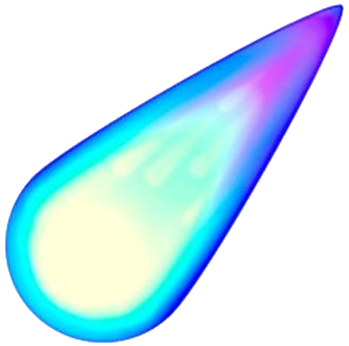

COMET

What are the implications of this knowledge representation?

Commonsense Knowledge for any Situation

transformer-style architecture — input format is natural language

- event can be fully parsed

Kai knew that things were getting out of control and managed to keep his temper in check

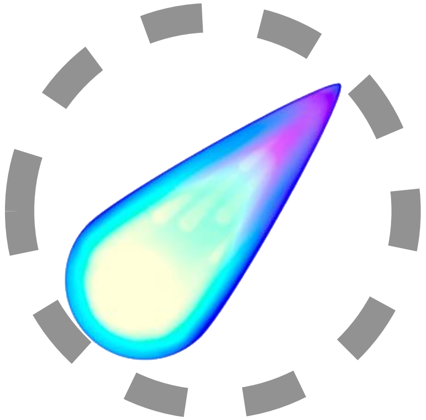

Commonsense Knowledge for any Situation

transformer-style architecture — input format is natural language

- event can be fully parsed
- knowledge generated **dynamically** from neural knowledge model

Kai knew that things were getting out of control and managed to keep his temper in check

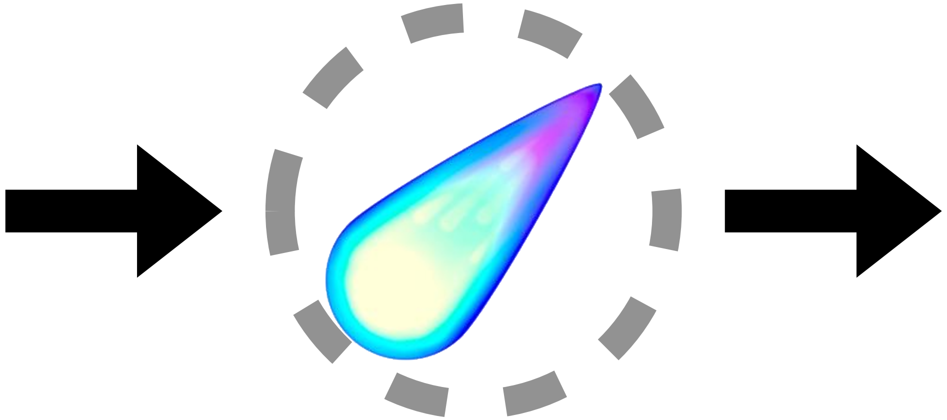

Kai wants to avoid trouble

Kai intends to be calm

Kai stays calm

Kai is viewed as cautious

Ways of combining them

During training

- Such as in reinforcement learning or retrieval-augmented generation (RAG)

After training

- Like a symbolic “wrapper” – helps validate what the NN is doing

Others??

VerbNet v3.4

https://verbs.colorado.edu/verbnet/

Verb classes based on Beth Levin (1993)

Data Source: hand-crafted

Languages: English

Use: raw data or my code

Demo: https://uvi.colorado.edu/uvi_search

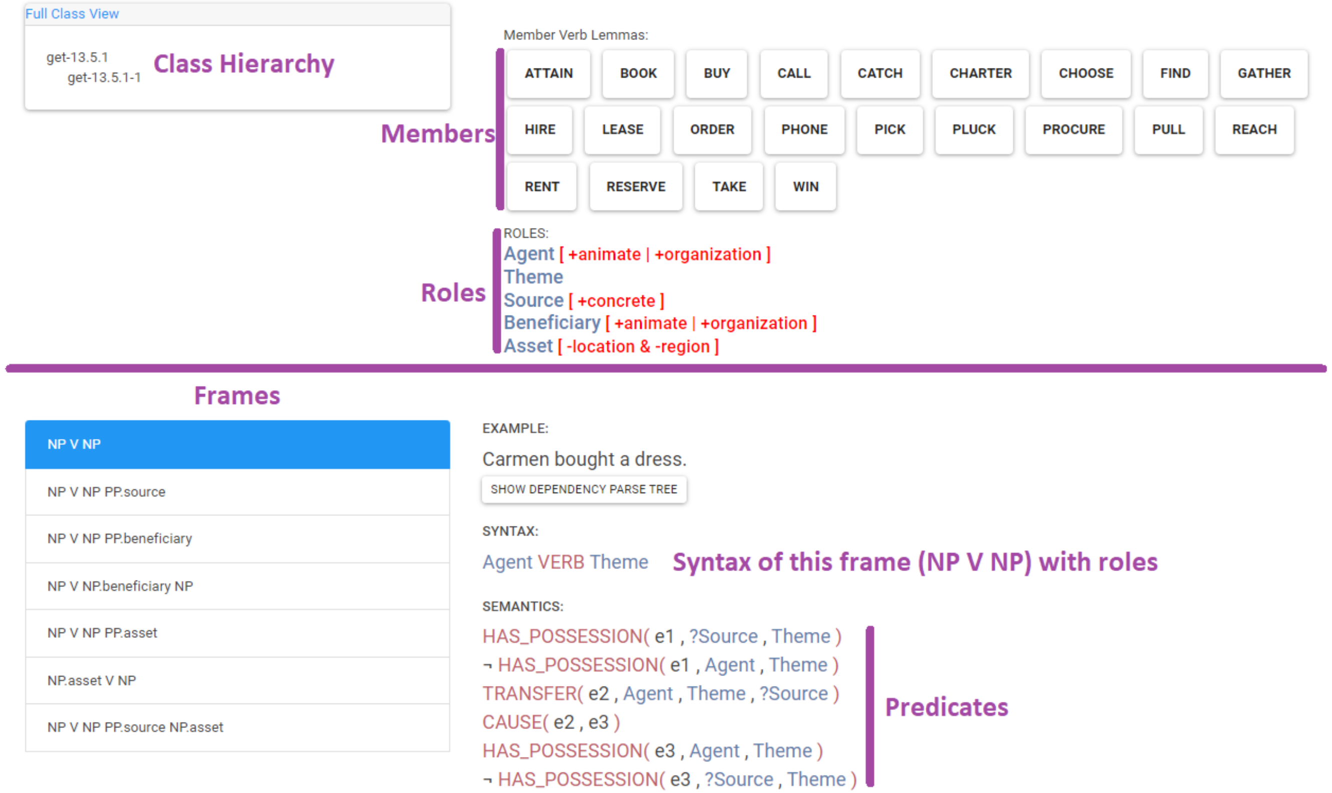

K. Kipper Schuler, "VerbNet: A Broad-Coverage, Comprehensive Verb Lexicon," University of Pennsylvania, 2005.
Levin, B. (1993) "English Verb Classes and Alternations: A Preliminary Investigation", University of Chicago Press, Chicago, IL.

Using VerbNet

Jen sent the book to Remy from Baltimore.

ROLES Agent Theme Destination Initial_Location

has_location(e1, book, Baltimore)
do(e2, **Jen**)
cause(e2, e3)
motion(e3, book)
!has_location(e3, book, Baltimore)
has_location(e4, book, Remy)

PREDICATES

Initial_Location : location

Theme : concrete

Agent : animate or organization

SELECTIONAL RESTRICTIONS

Pre-Conditions and Effects

Jen sent the book to Remy from Baltimore.

Pre-Conditions

has_location(e1, book, Baltimore)
do(e2, **Jen**)

cause(e2, e3)

motion(e3, book)

!has_location(e3, book, Baltimore)

has_location(e4, book, Remy)

Effects

~~Baltimore : location~~

book : concrete

Jen : animate or organization

Pre-Conditions and Effects

Jen sent the book to Remy from Baltimore.

Pre-Conditions	Effects
has_location(e1, book, Baltimore)	~~do(e2, **Jen**)~~
Baltimore : location	~~cause(e2, e3)~~
book : concrete	~~motion(e3, book)~~
Jen : animate or organization	!has_location(e3, book, Baltimore)
	has_location(e4, book, Remy)

Pre-Conditions and Effects

Jen sent the book to Remy from Baltimore.

Pre-Conditions → Event → Effects

has_location(book, Baltimore)
Baltimore : location
book : concrete
Jen : animate or organization

!has_location(book, Baltimore)
has_location(book, Remy)

L. J. Martin, "Neurosymbolic Automated Story Generation," PhD, Georgia Institute of Technology, Atlanta, GA, 2021. https://smartech.gatech.edu/handle/1853/64643

Resulting State Representation

Jen sent the book to Remy from Baltimore.

Baltimore : location
book : concrete
Jen : animate or organization
!has_location(book, Baltimore)
has_location(book, Remy)

How does a neural network fit in here?

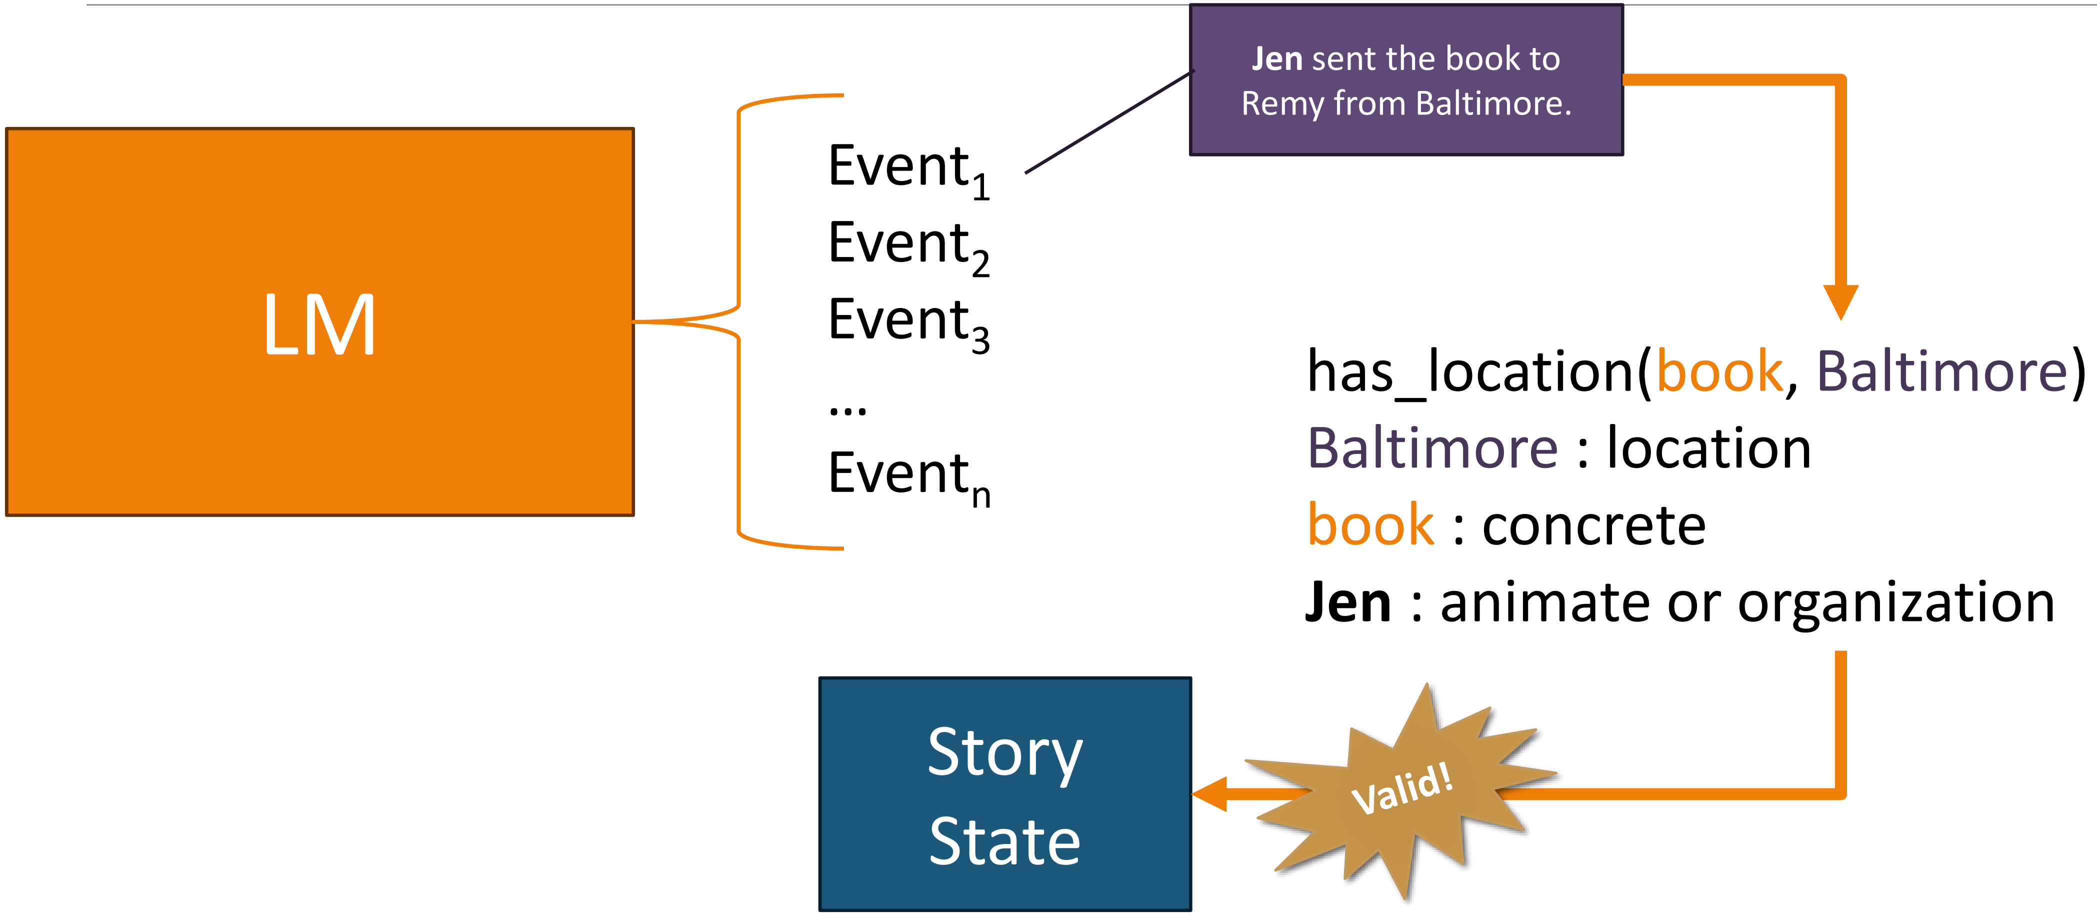

Knowledge Check

1. Why might neurosymbolic systems still be useful with today's few-shot LMs?
2. What are some ways you would integrate a knowledge base into a modern LM?