Commonsense AI: Myth and Truth

— 🍔 *cheeseburger stabbing redux, 2021 edition* —

Yejin Choi

Paul G. Allen School of Computer Science & Engineering
University of Washington &
Allen Institute for Artificial Intelligence


Commonsense AI: Myth and Truth
— 🍔 cheeseburger stabbing redux, 2021 edition —

Year 2020

ACL Commonsense Tutorial

— 2nd most popular (among 8 tutorials) —

https://homes.cs.washington.edu/~msap/acl2020-commonsense/

5k = # of the main conference registration
1.3k = # of our tutorial registration
3.4k = # of view counts on our recorded lectures

Vered Shwartz

Maarten Sap

Antoine Bosselut

Dan Roth

Yejin Choi

Circa 2017…

That’s a research topic of 70s & 80s

😱

Shouldn’t work on it since it won’t work

😱

Commonsense ?

Don’t even say that word

😱

2017 — 2020

What is commonsense?

- It’s what everyone knows and agrees on?
- language is irrelevant to commonsense?
- Too hard to define precisely, thus shouldn’t work on it

Should we or should we not?

- Commonsense AI is an impossible goal (ever)
- That’s a research topic of 70s and 80s
- Maybe only possible in the faraway future

Commonly held beliefs

Truth or Myth ?

- Knowledge and reasoning are distinct and exclusive
- Language is in the way of reasoning; let's do formal logics
- Language is not symbols. Words and numbers are, but not language at large
- Humans acquire commonsense completely un-/self-supervised, thus so should machines

What is commonsense?

- It's what everyone knows and agrees on?
- language is irrelevant to commonsense?
- Too hard to define precisely, thus shouldn't work on it

Should we or should we not?

- Commonsense AI is an impossible goal (ever)
- That's a research topic of 70s and 80s
- Maybe only possible in the faraway future

The Curious Case of **Cheeseburger Stabbing**

*— An example repeatedly appeared my talks between **Mar 2017** and **May 2018** —*

The Curious Case of 🍔 "**Cheeseburger Stabbing**"

Someone stabbed a cheeseburger?
A cheeseburger stabbed someone?
A cheeseburger stabbed another cheeseburger?
Someone stabbed someone else over a cheeseburger?

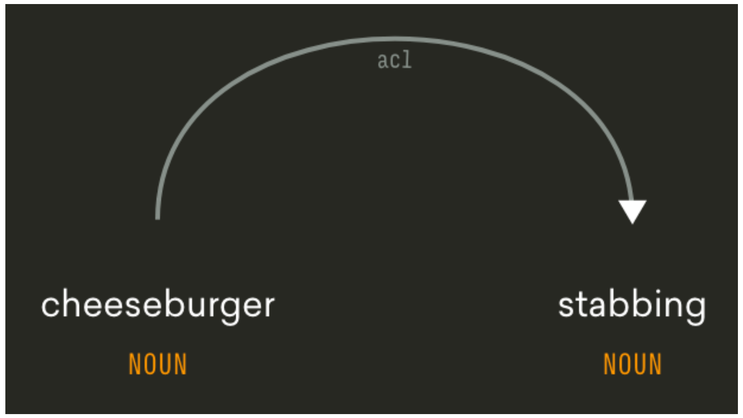

The Curious Case of 🍔 "**Cheeseburger Stabbing**"

Physical Commonsense: not possible to stab using a burger

Social Commonsense: stabbing someone is bad

what is said +

what is *not* said

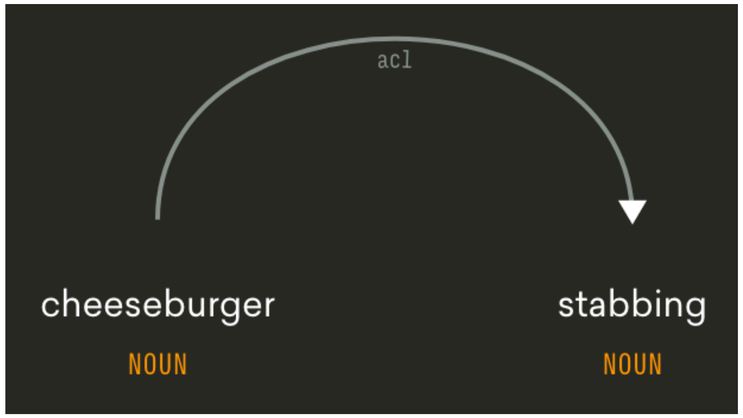

Someone stabbed a cheeseburger?

A cheeseburger stabbed someone?

A cheeseburger stabbed another cheeseburger?

Someone stabbed **someone** else over a cheeseburger?

The Curious Case of 🍔"Cheeseburger Stabbing"

Vered Shwartz
@VeredShwartz

plying to @shacharmirkin and @AllThingsLing

Thanks! I first heard it last year in @YejinChoinka 's talk, and I used it as a difficult example to test my noun

cheeseburger" 😂

12:10 AM · Aug 31, 2018 · Twitter for Android

2018:
"Stabbing is a crime punishable by cheeseburger" 😂

2020:
"stabbing of a cheeseburger" per GPT-3

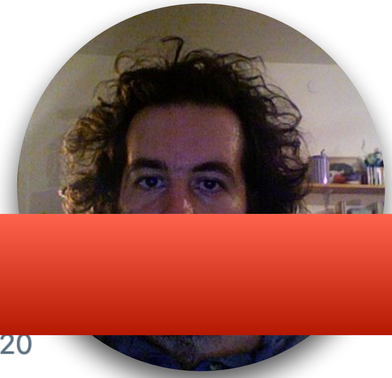

2021: Stay tuned for the 2021 edition!

((((yoav' ()ↄ)()ↄ)))) @yoavgo · Jul 18, 2020

Q: parsley cake
A: cake made of parsley

Q: cheeseburger stabbing
A: stabbing of a cheeseburger

Path to commonsense?

Brute force larger networks with deeper layers?

You don’t reach to the moon
by making the tallest building in the world
taller

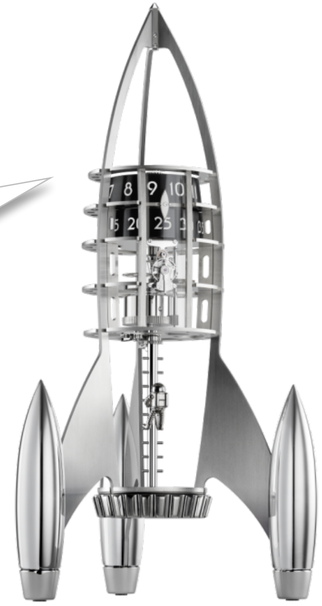

Path to commonsense?

Obligatory

🔥Controversial Remarks🔥 of the Day

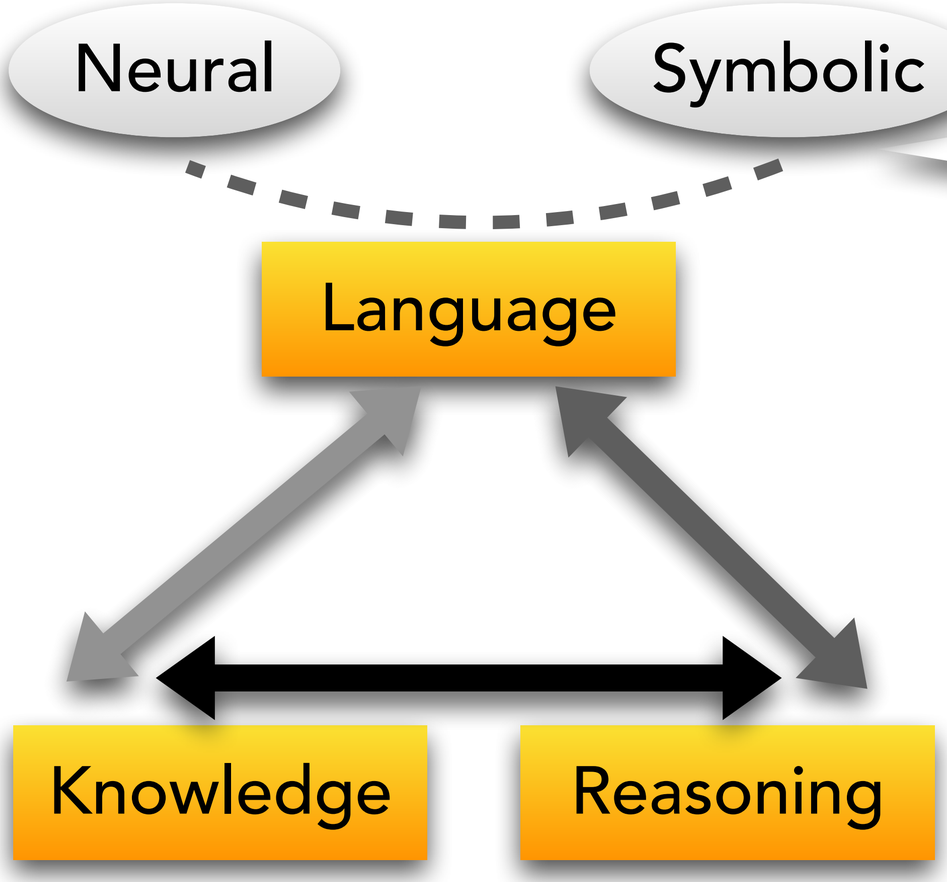

1. the continuum between **knowledge** and **reasoning**
2. the interplay between **reasoning** and **language generation**
3. the blend between **neural** vs **symbolic** representation

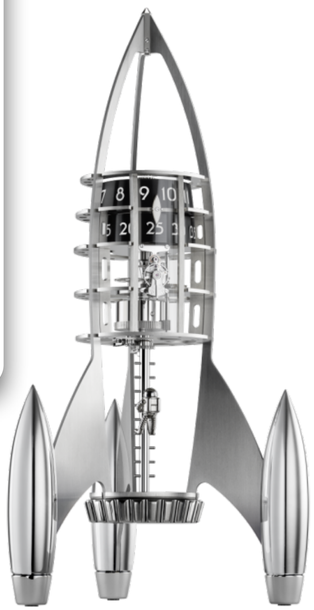

Language

Reasoning

Induction — from specific to general

Deduction — from general to specific

Abduction — "*why* something happened"

(or reasoning about the the probable explanation)

Peirce 1965

Counterfactual — "*what if* something else happened"

Back to the Future:

Unsupervised Backprop-based Decoding for Counterfactual and Abductive Commonsense Reasoning

EMNLP 2020

Lianhui Qin

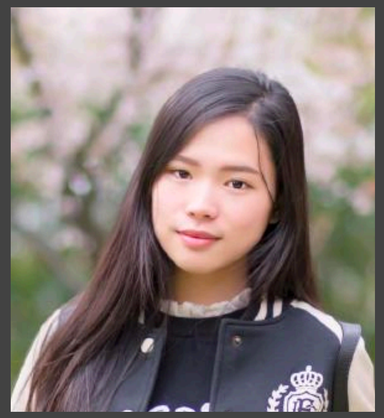

Vered Shwartz

Peter West

Chandra Bhagavatula

Jena Hwang

Ronan LeBras

Antoine Bosselut

Me

Abductive Reasoning

(Bhagavatula et al., 2019)

Past Observation

Ray hung a tire on a rope to make his daughter a swing.

What happened in between?

Future Observation

Ray ran to his daughter to make sure she was okay.

Abductive Reasoning

(Bhagavatula et al., 2019)

Past Observation

Ray hung a tire on a rope to make his daughter a swing.

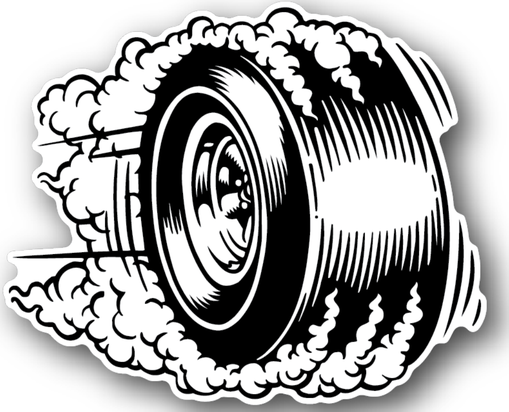

Hypothesis

She hit the rope and the tire fell on top of her.

Future Observation

Ray ran to his daughter to make sure she was okay.

Abduction — "*why* something happened"

(or reasoning about the the probable explanation)

Peirce 1965

Abductive Reasoning

(Bhagavatula et al., 2019)

Past Observation

Ray hung a tire on a rope to make his daughter a swing.

Hypothesis

She hit the rope and the tire fell on top of her.

Future Observation

Ray ran to his daughter to make sure she was okay.

Counterfactual Reasoning

(Qin et al., 2019)

An example from the
“TimeTravel” dataset
(Qin et al., EMNLP 2019)

Zeke was throwing a party.

All his friends were dressing up for this Halloween party. Story context changes…

Zeke thought about being a vampire or a wizard.

Then he decided on a scarier costume.

Zeke dressed up like a skeleton.

What if this is a Game of Thrones themed party instead of a Halloween party? 🤔

An example from the "TimeTravel" dataset (Qin et al., EMNLP 2019)

Zeke was throwing a party.

All his friends were dressing up for this Halloween party. Story context changes…

Zeke thought about being a vampire or a wizard.

Then he decided on a scarier costume.

Zeke dressed up like a skeleton.

Story ending doesn't make sense now…

What if this is a Game of Thrones themed party instead of a Halloween party? 🤔

An example from the "TimeTravel" dataset (Qin et al., EMNLP 2019)

Zeke was throwing a party.

All his friends were dressing up for this Halloween party.

Story context changes…

Zeke thought about being a vampire or a wizard.

Then he decided on a scarier costume.

Zeke dressed up like a skeleton.

Story ending doesn't make sense now…

What if this is a Game of Thrones themed party instead of a Halloween party?

Counterfactual Reasoning

(Qin et al., 2019)

Reasoning about the alternative ***future*** based on counterfactual ***past***.

An example from the "TimeTravel" dataset (Qin et al., EMNLP 2019)

Zeke was throwing a party.

All his friends were dressing up for this Halloween party. Story context changes…

Zeke thought about ~~being a vampire or a wizard.~~

~~Then he decided on a scarier costume.~~

Zeke dressed up like a ~~skeleton.~~

Story ending doesn't make sense now…

What if this is a Game of Thrones themed party instead of a Halloween party?

Zeke was throwing a party.

All his friends were dressing up for this Halloween party.

Zeke thought about *Lannister, but he didn't want to look like a Lannister.*

He wanted to look like a Stark.

Zeke dressed up like a *Stark*.

Only do minimal edit!

Consistent now!!

Abductive Reasoning

(Bhagavatula et al., 2019)

Past Observation

Ray hung a tire on a rope to make his daughter a swing.

Hypothesis

She hit the rope and the tire fell on top of her.

Future Observation

Ray ran to his daughter to make sure she was okay.

Story Context

Zeke was throwing a party.

All his friends were dressing up for this Halloween party.

All his friends were dressing up for this Game of Thrones themed party.

Rewritten Ending

Zeke thought about Lannister, but he didn't want to look like a Lannister.

He wanted to look like a Stark.

Zeke dressed up like a Stark.

Original Ending

Zeke thought about being a vampire or a wizard.

Then he decided on a scarier costume.

Zeke dressed up like a skeleton.

Counterfactual Reasoning

(Qin et al., 2019)

Abductive Reasoning

(Bhagavatula et al., 2019)

Both involve ***nonmonotonic reasoning*** with past context X and future constraint Z

Counterfactual Reasoning

(Qin et al., 2019)

Pretrained Language Models are successful on many tasks…

How are Pretrained LMs on the Nonmonotonic Reasoning?

Let’s first see the abductive case…

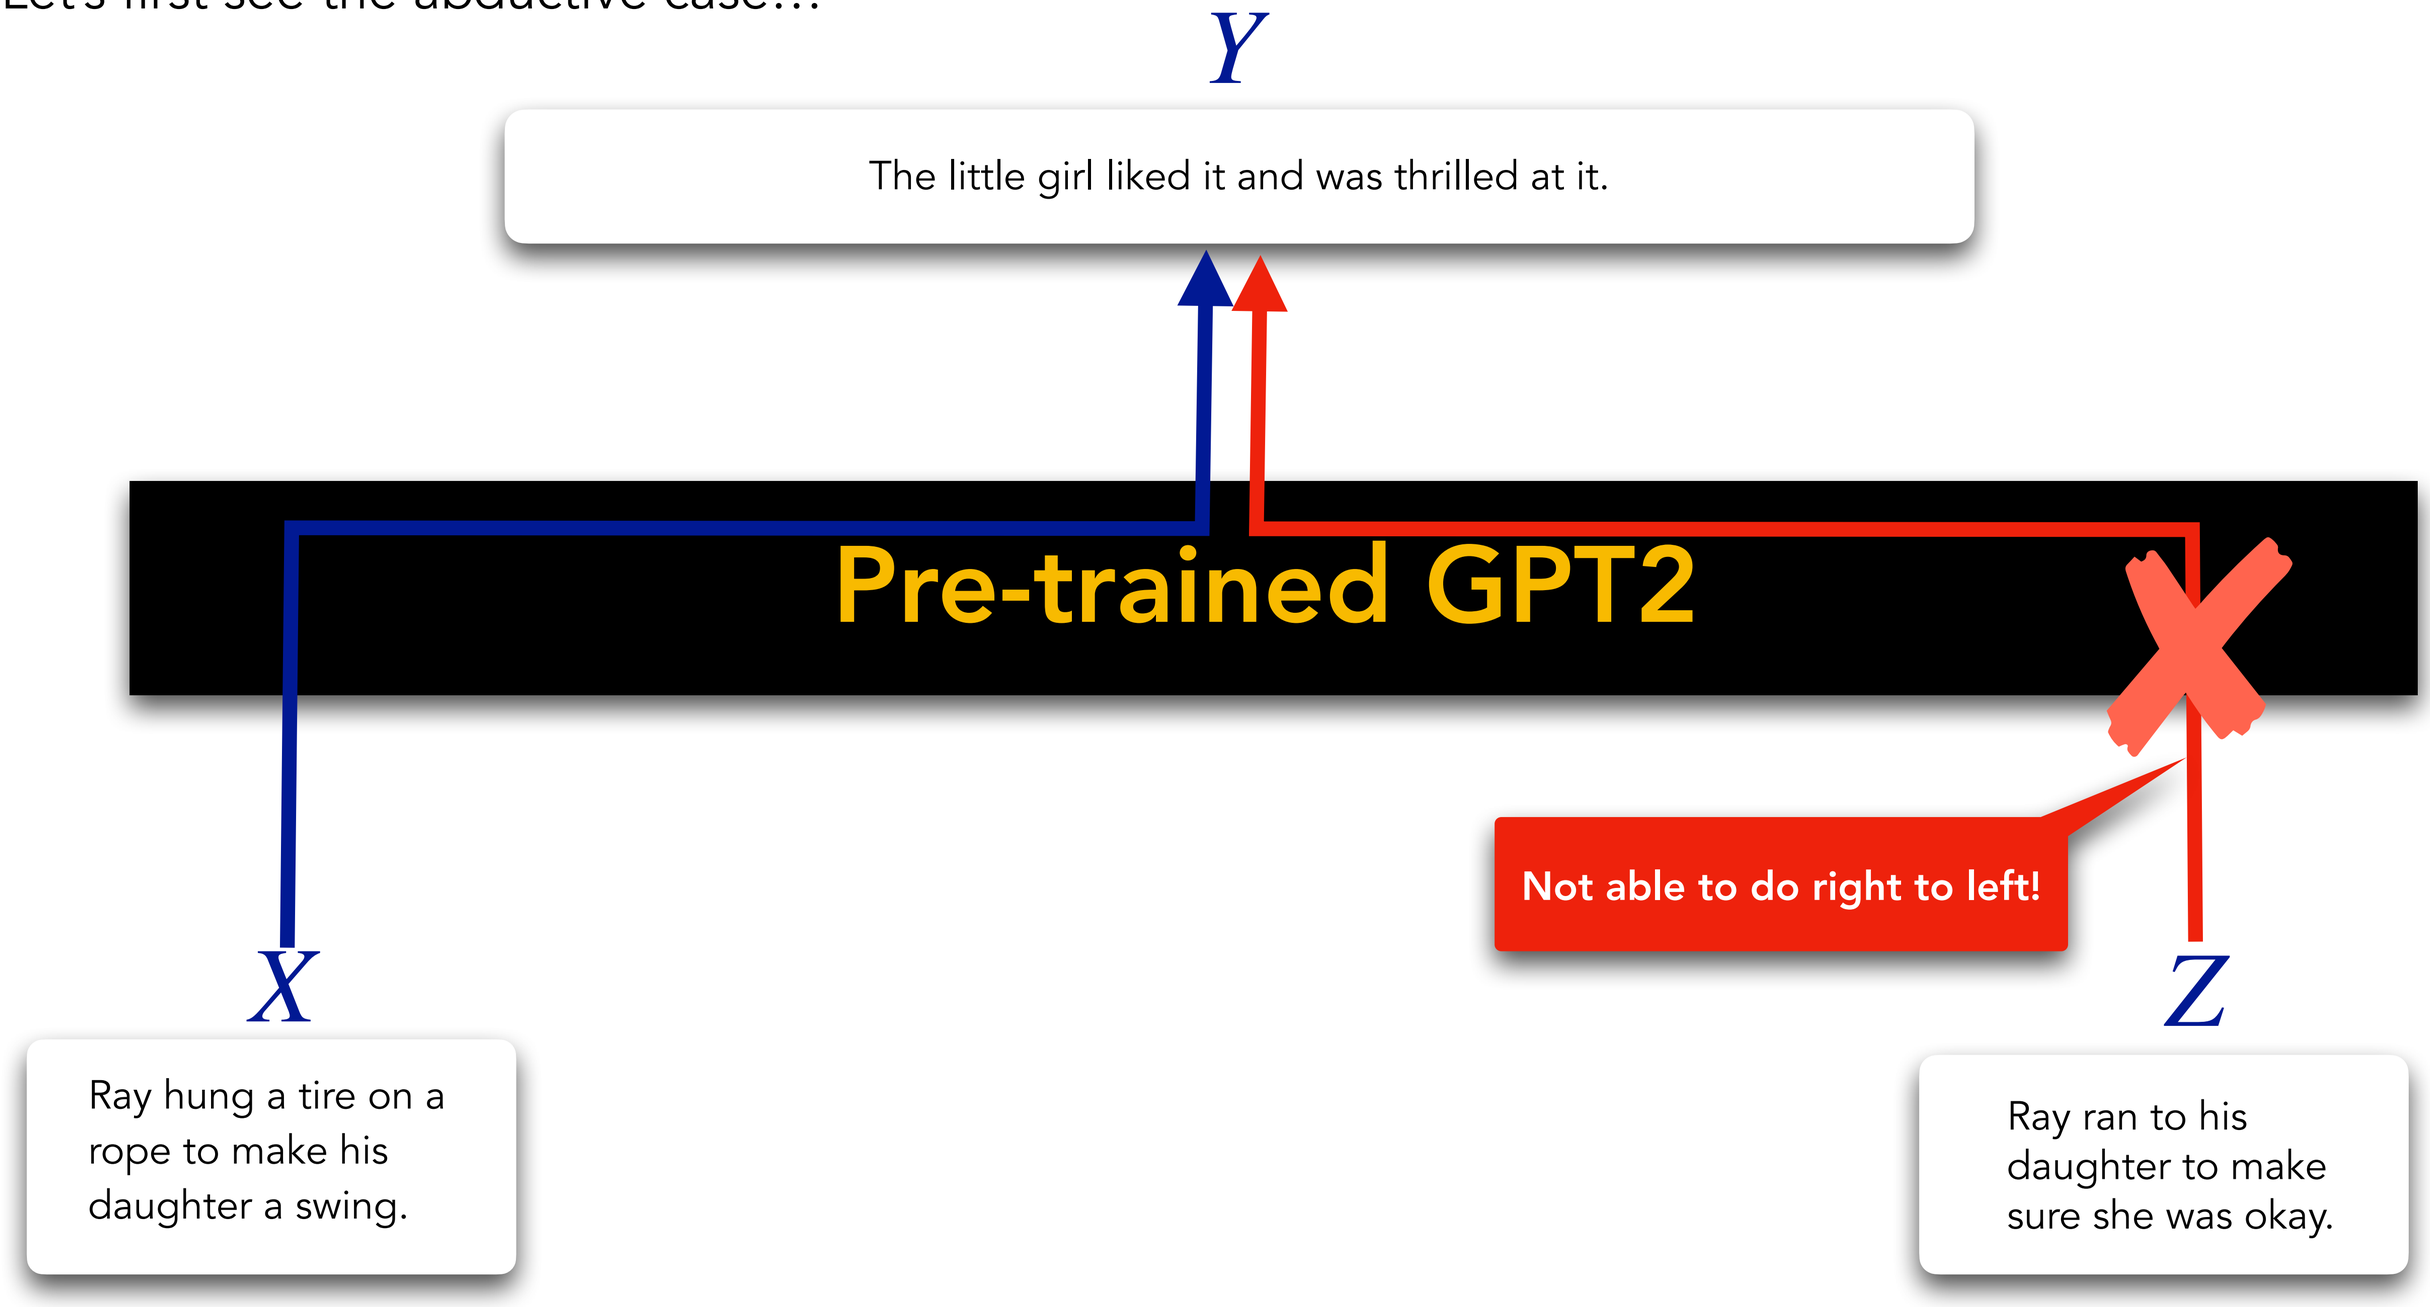

Why not just concatenate both direction?

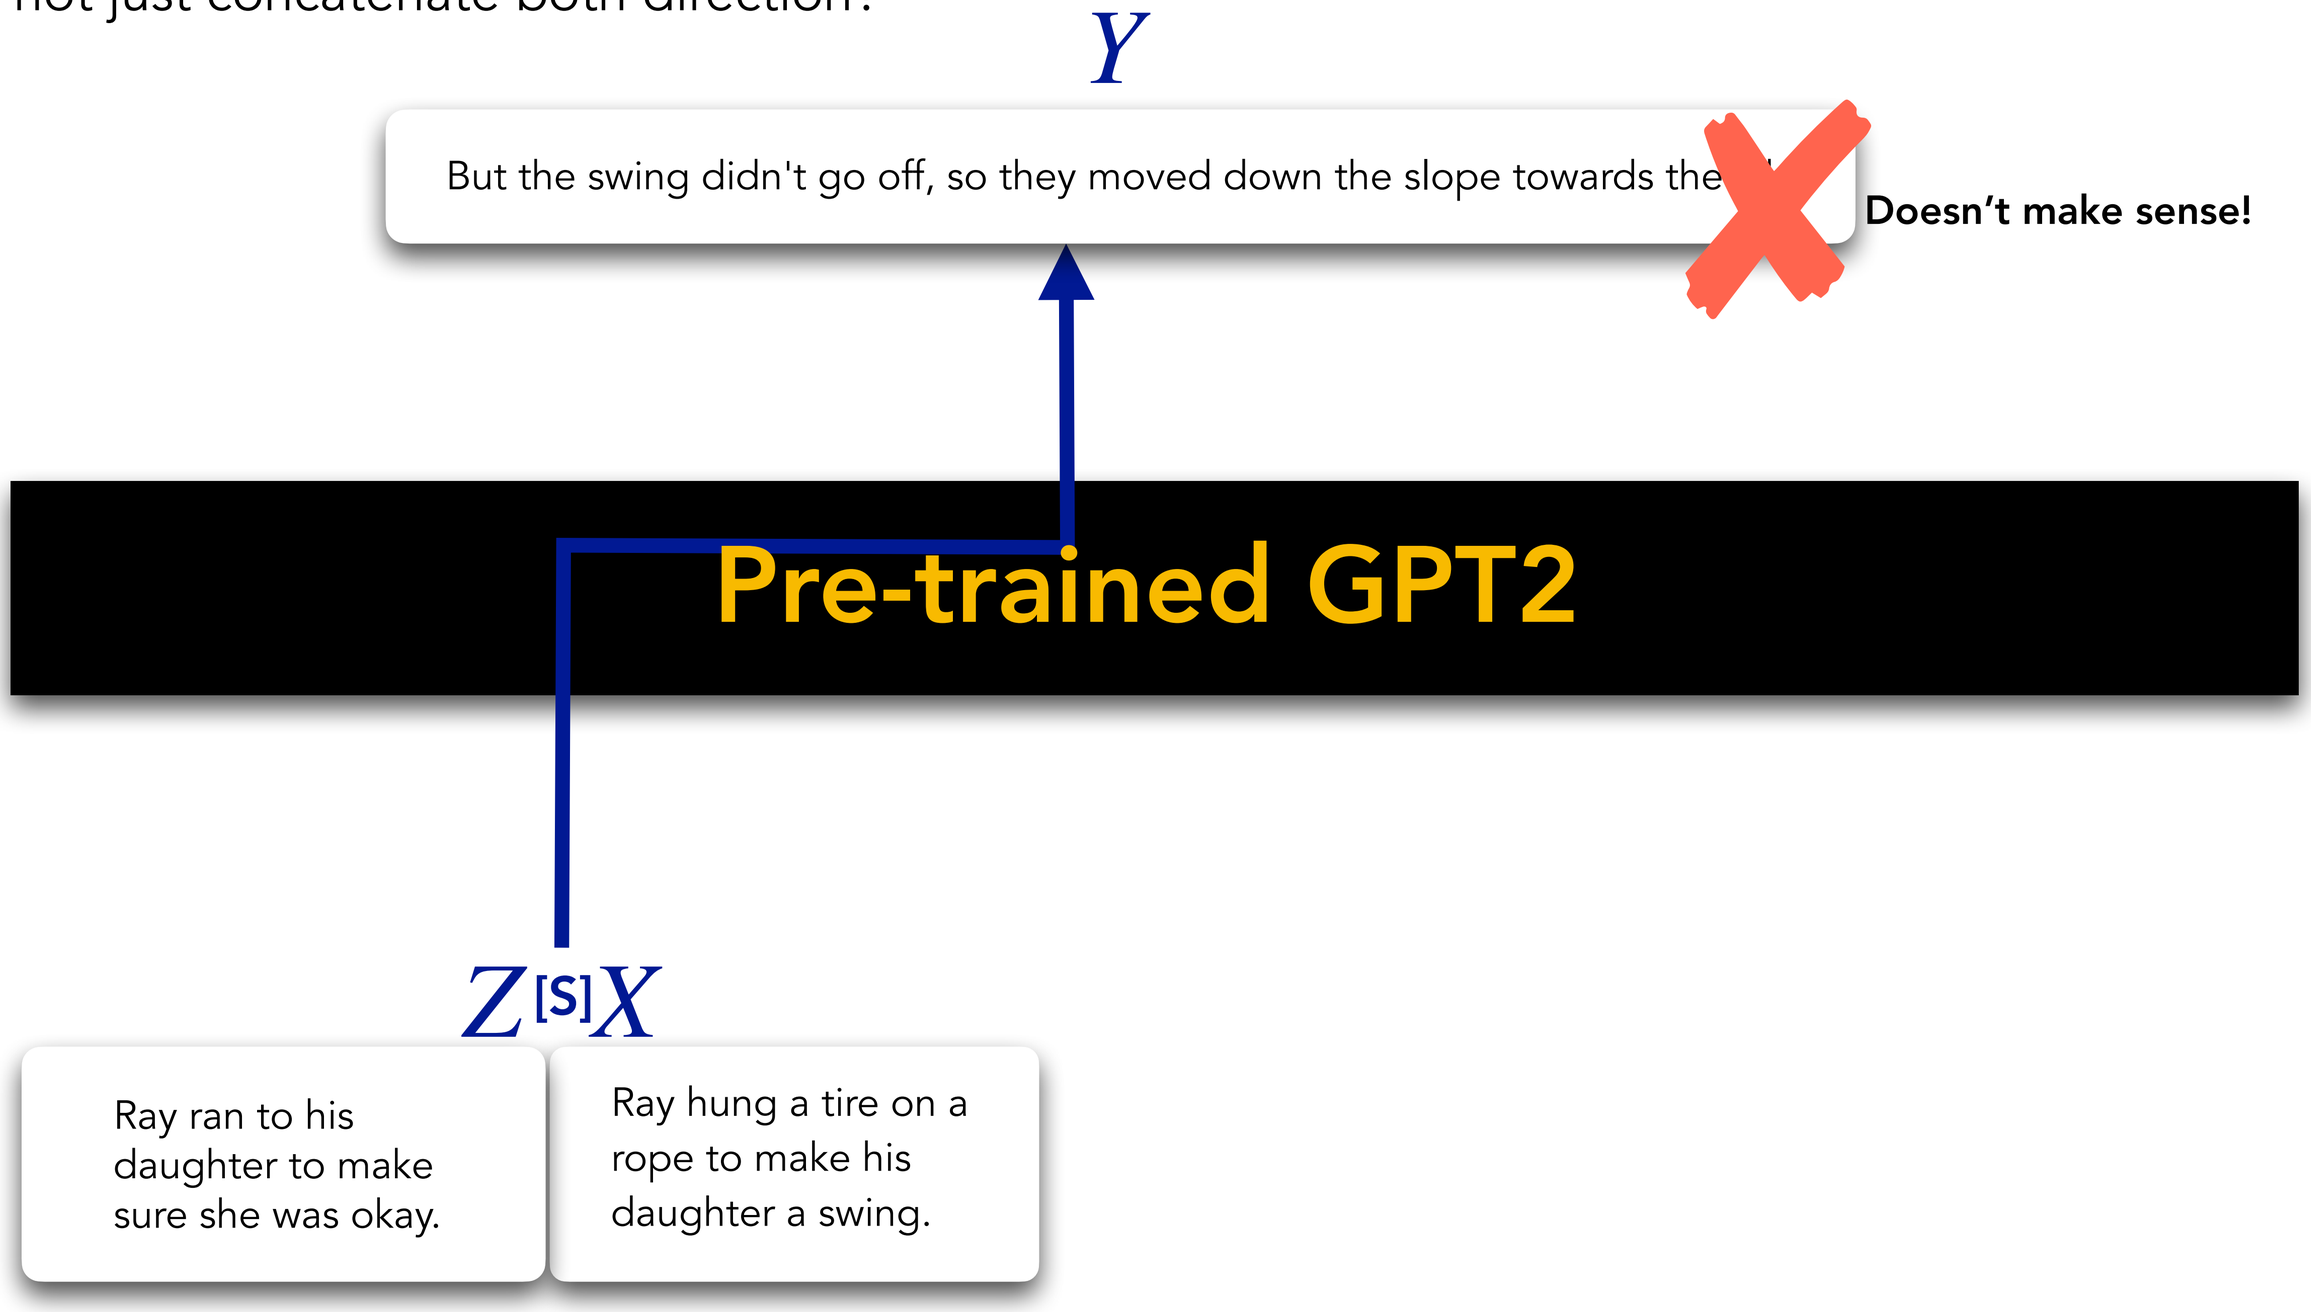

Try again?

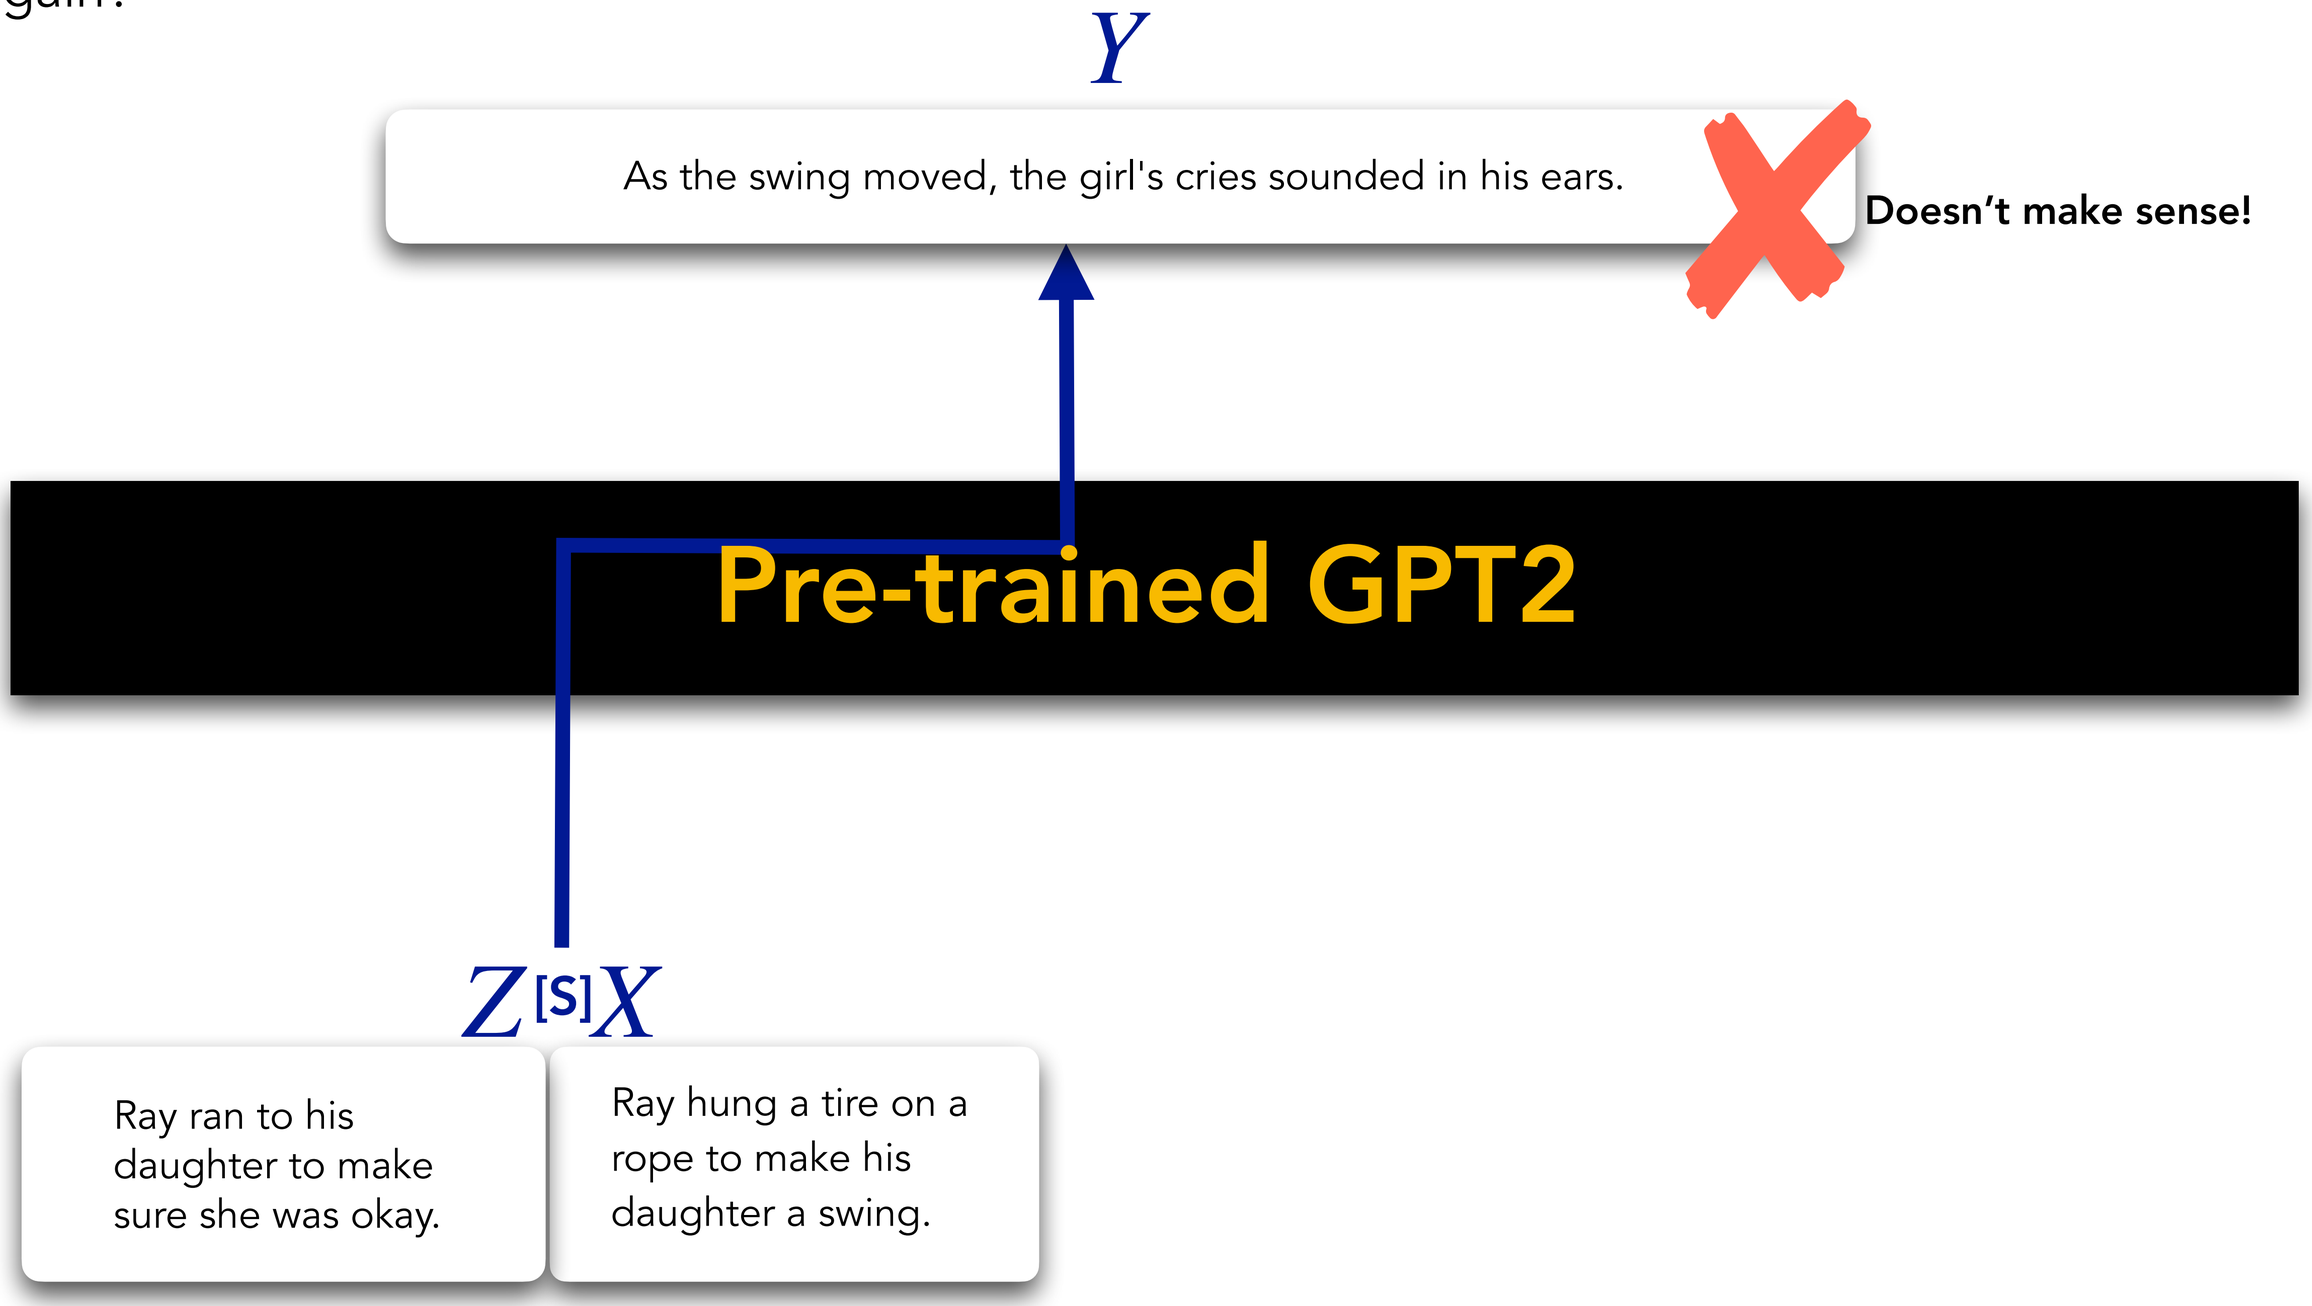
Y
As the swing moved, the girl's cries sounded in his ears.
Doesn't make sense!
Pre-trained GPT2
$Z^{[S]}X$
Ray ran to his daughter to make sure she was okay.
Ray hung a tire on a rope to make his daughter a swing.

Something might have been missing here…

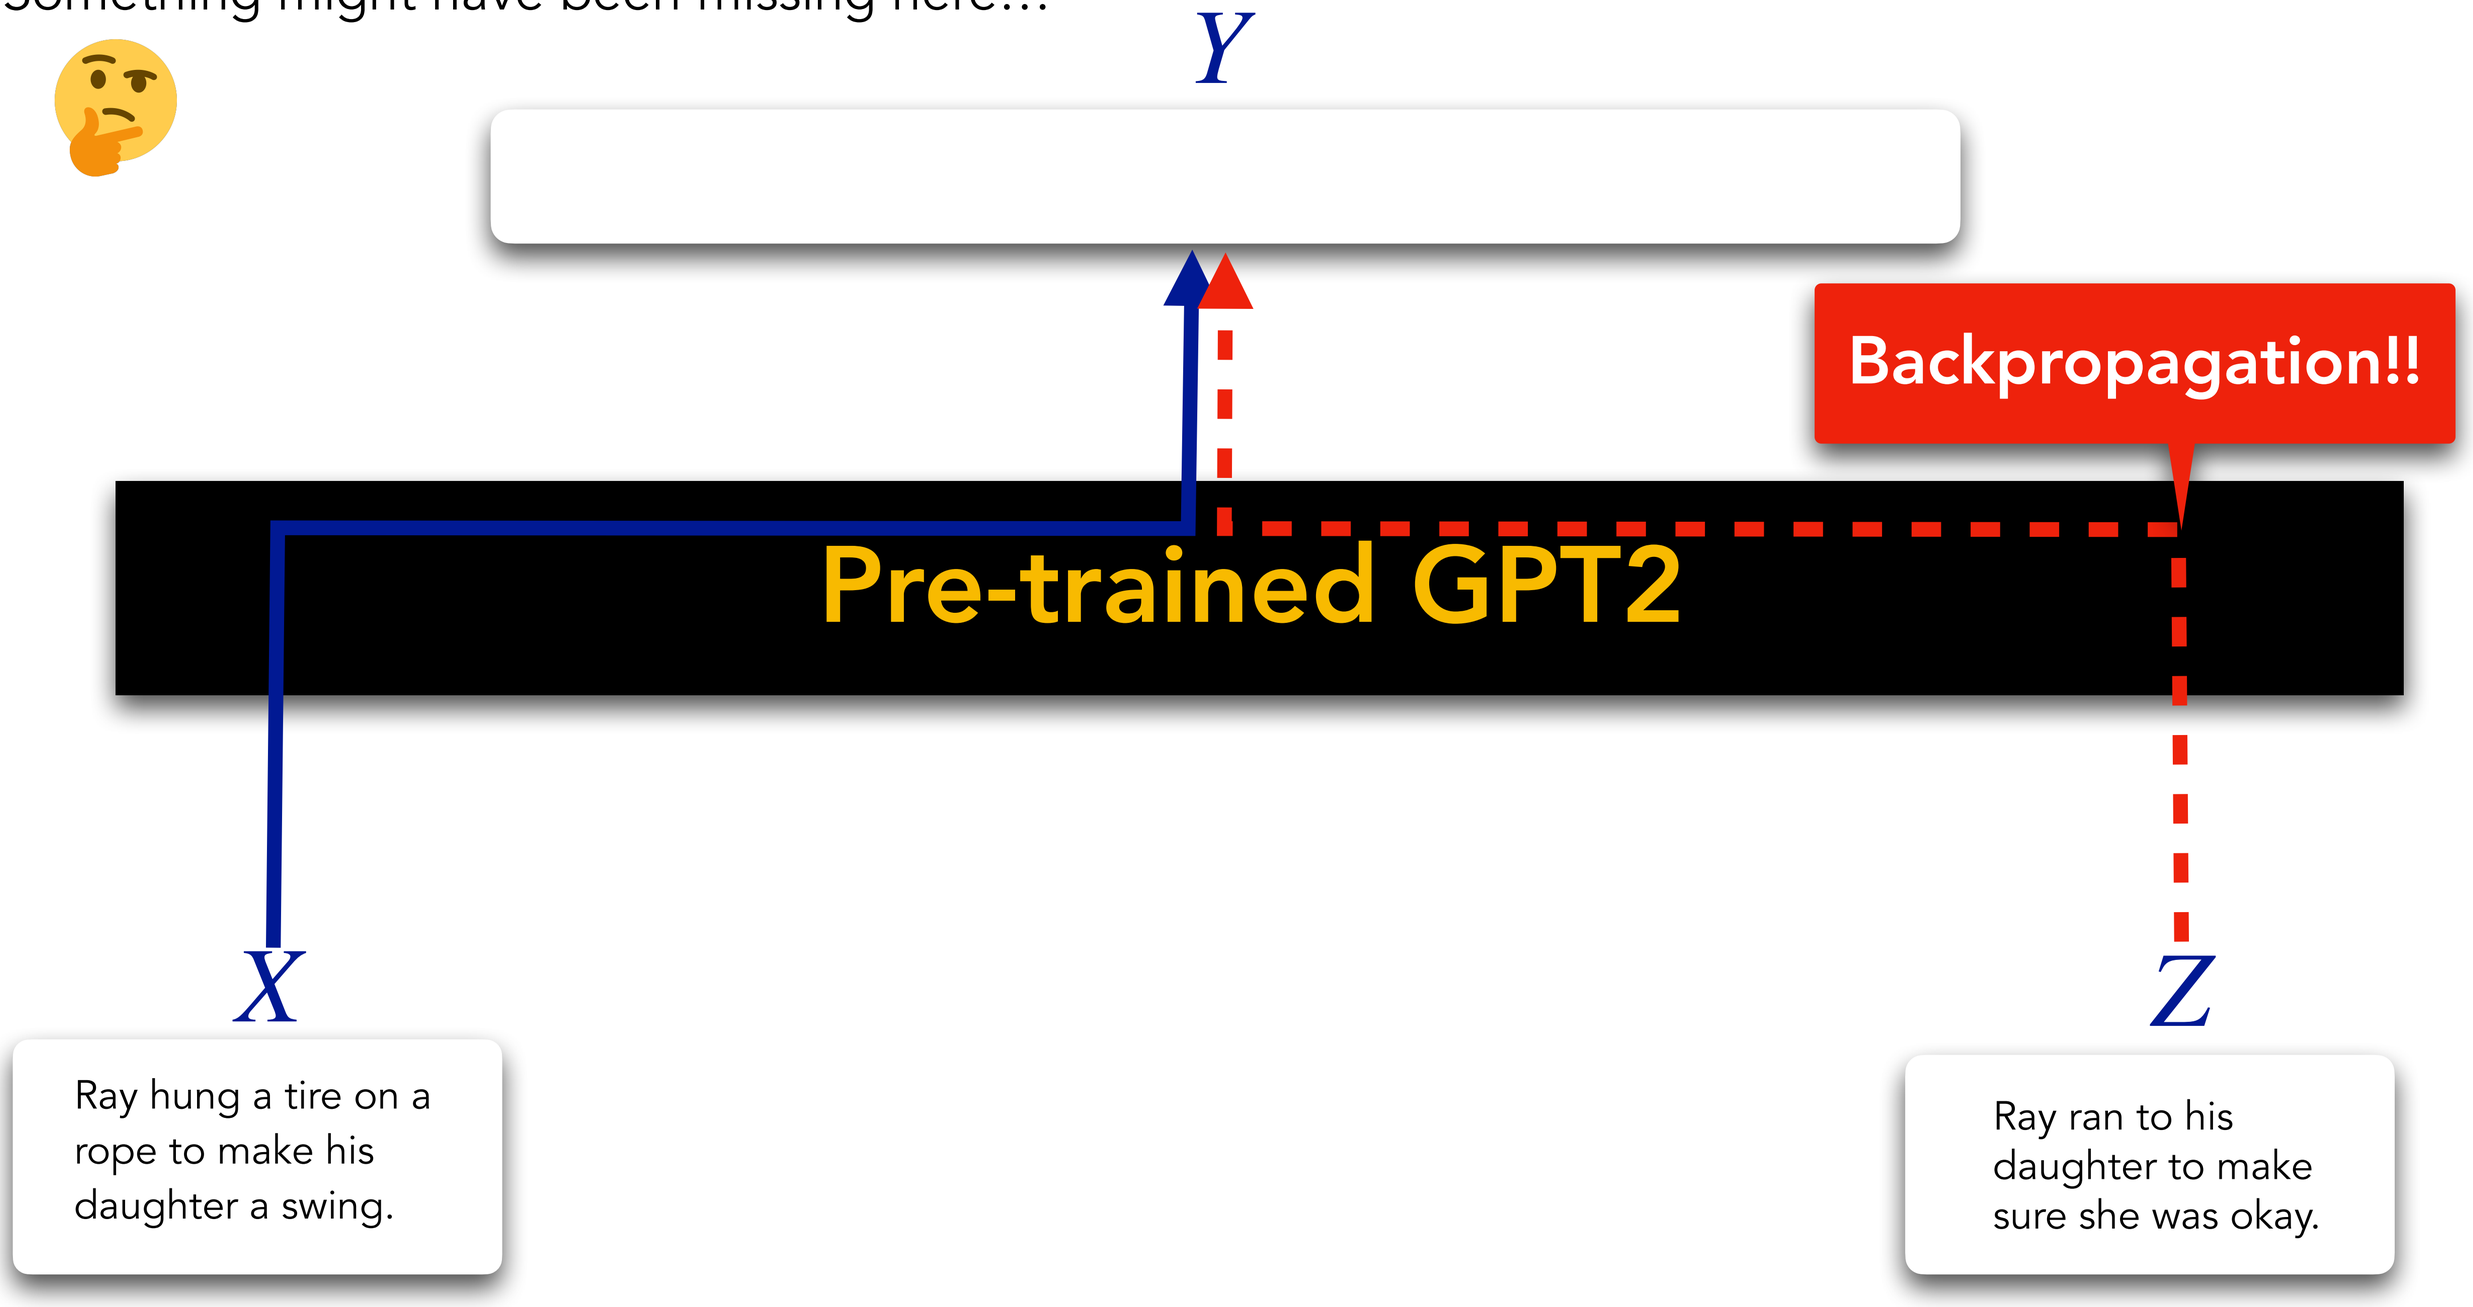

Inspired by “Image Style Transfer” (Gatys et al, 2016)…

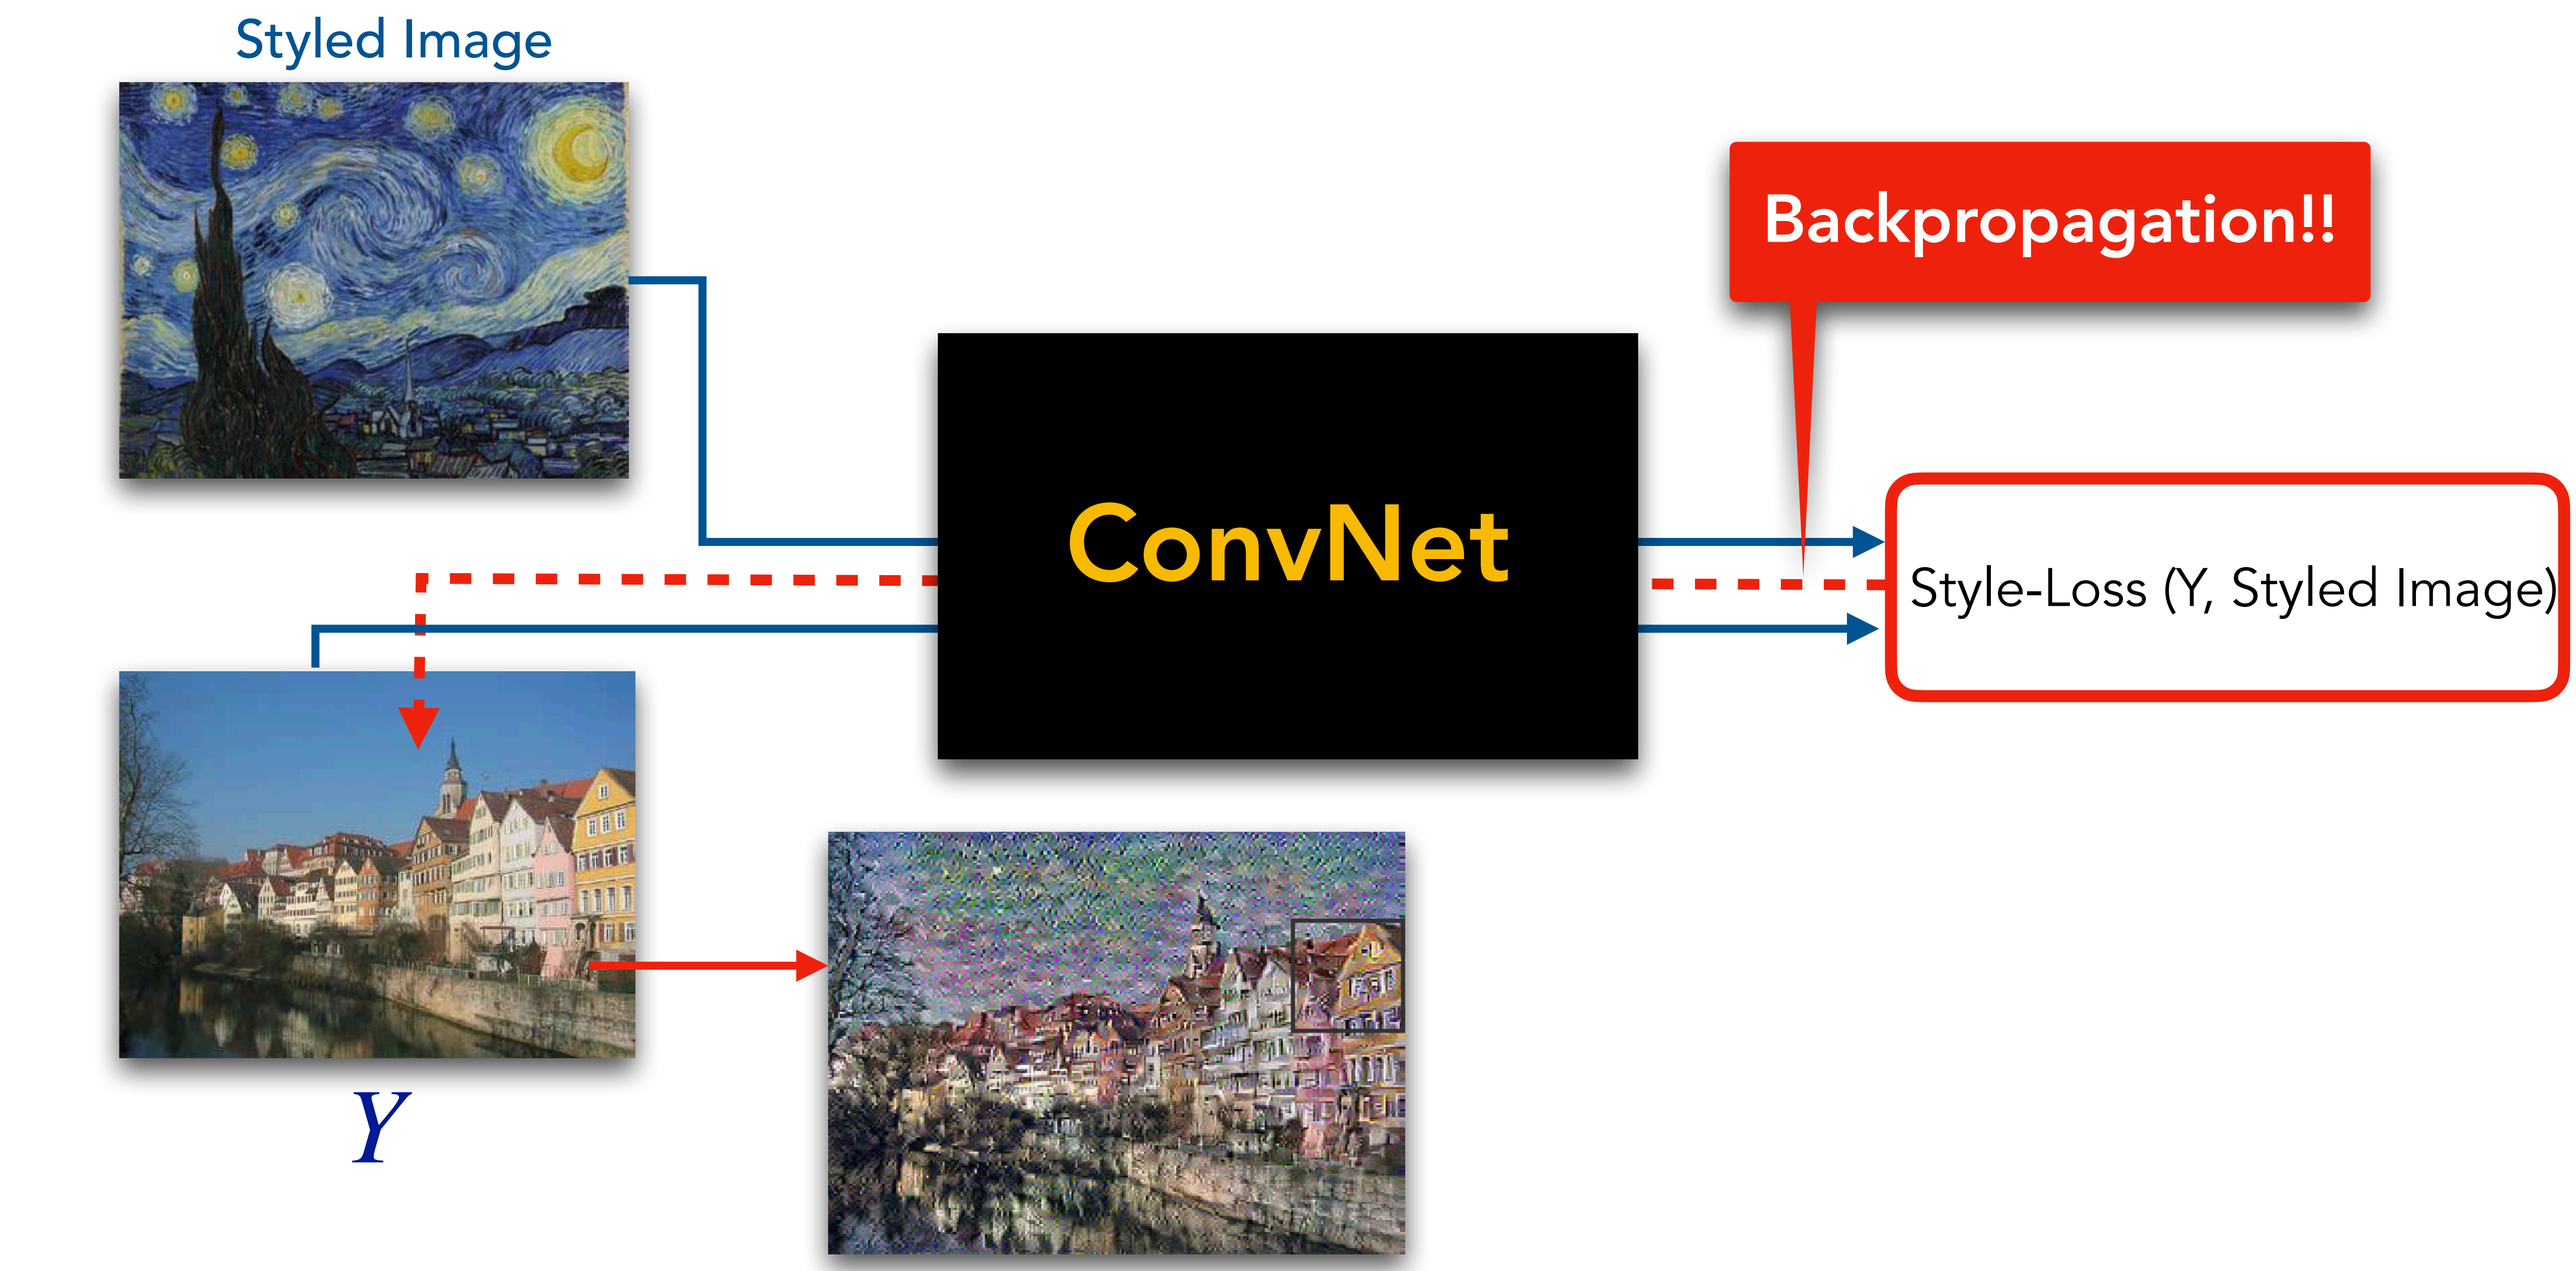

Back to our case …

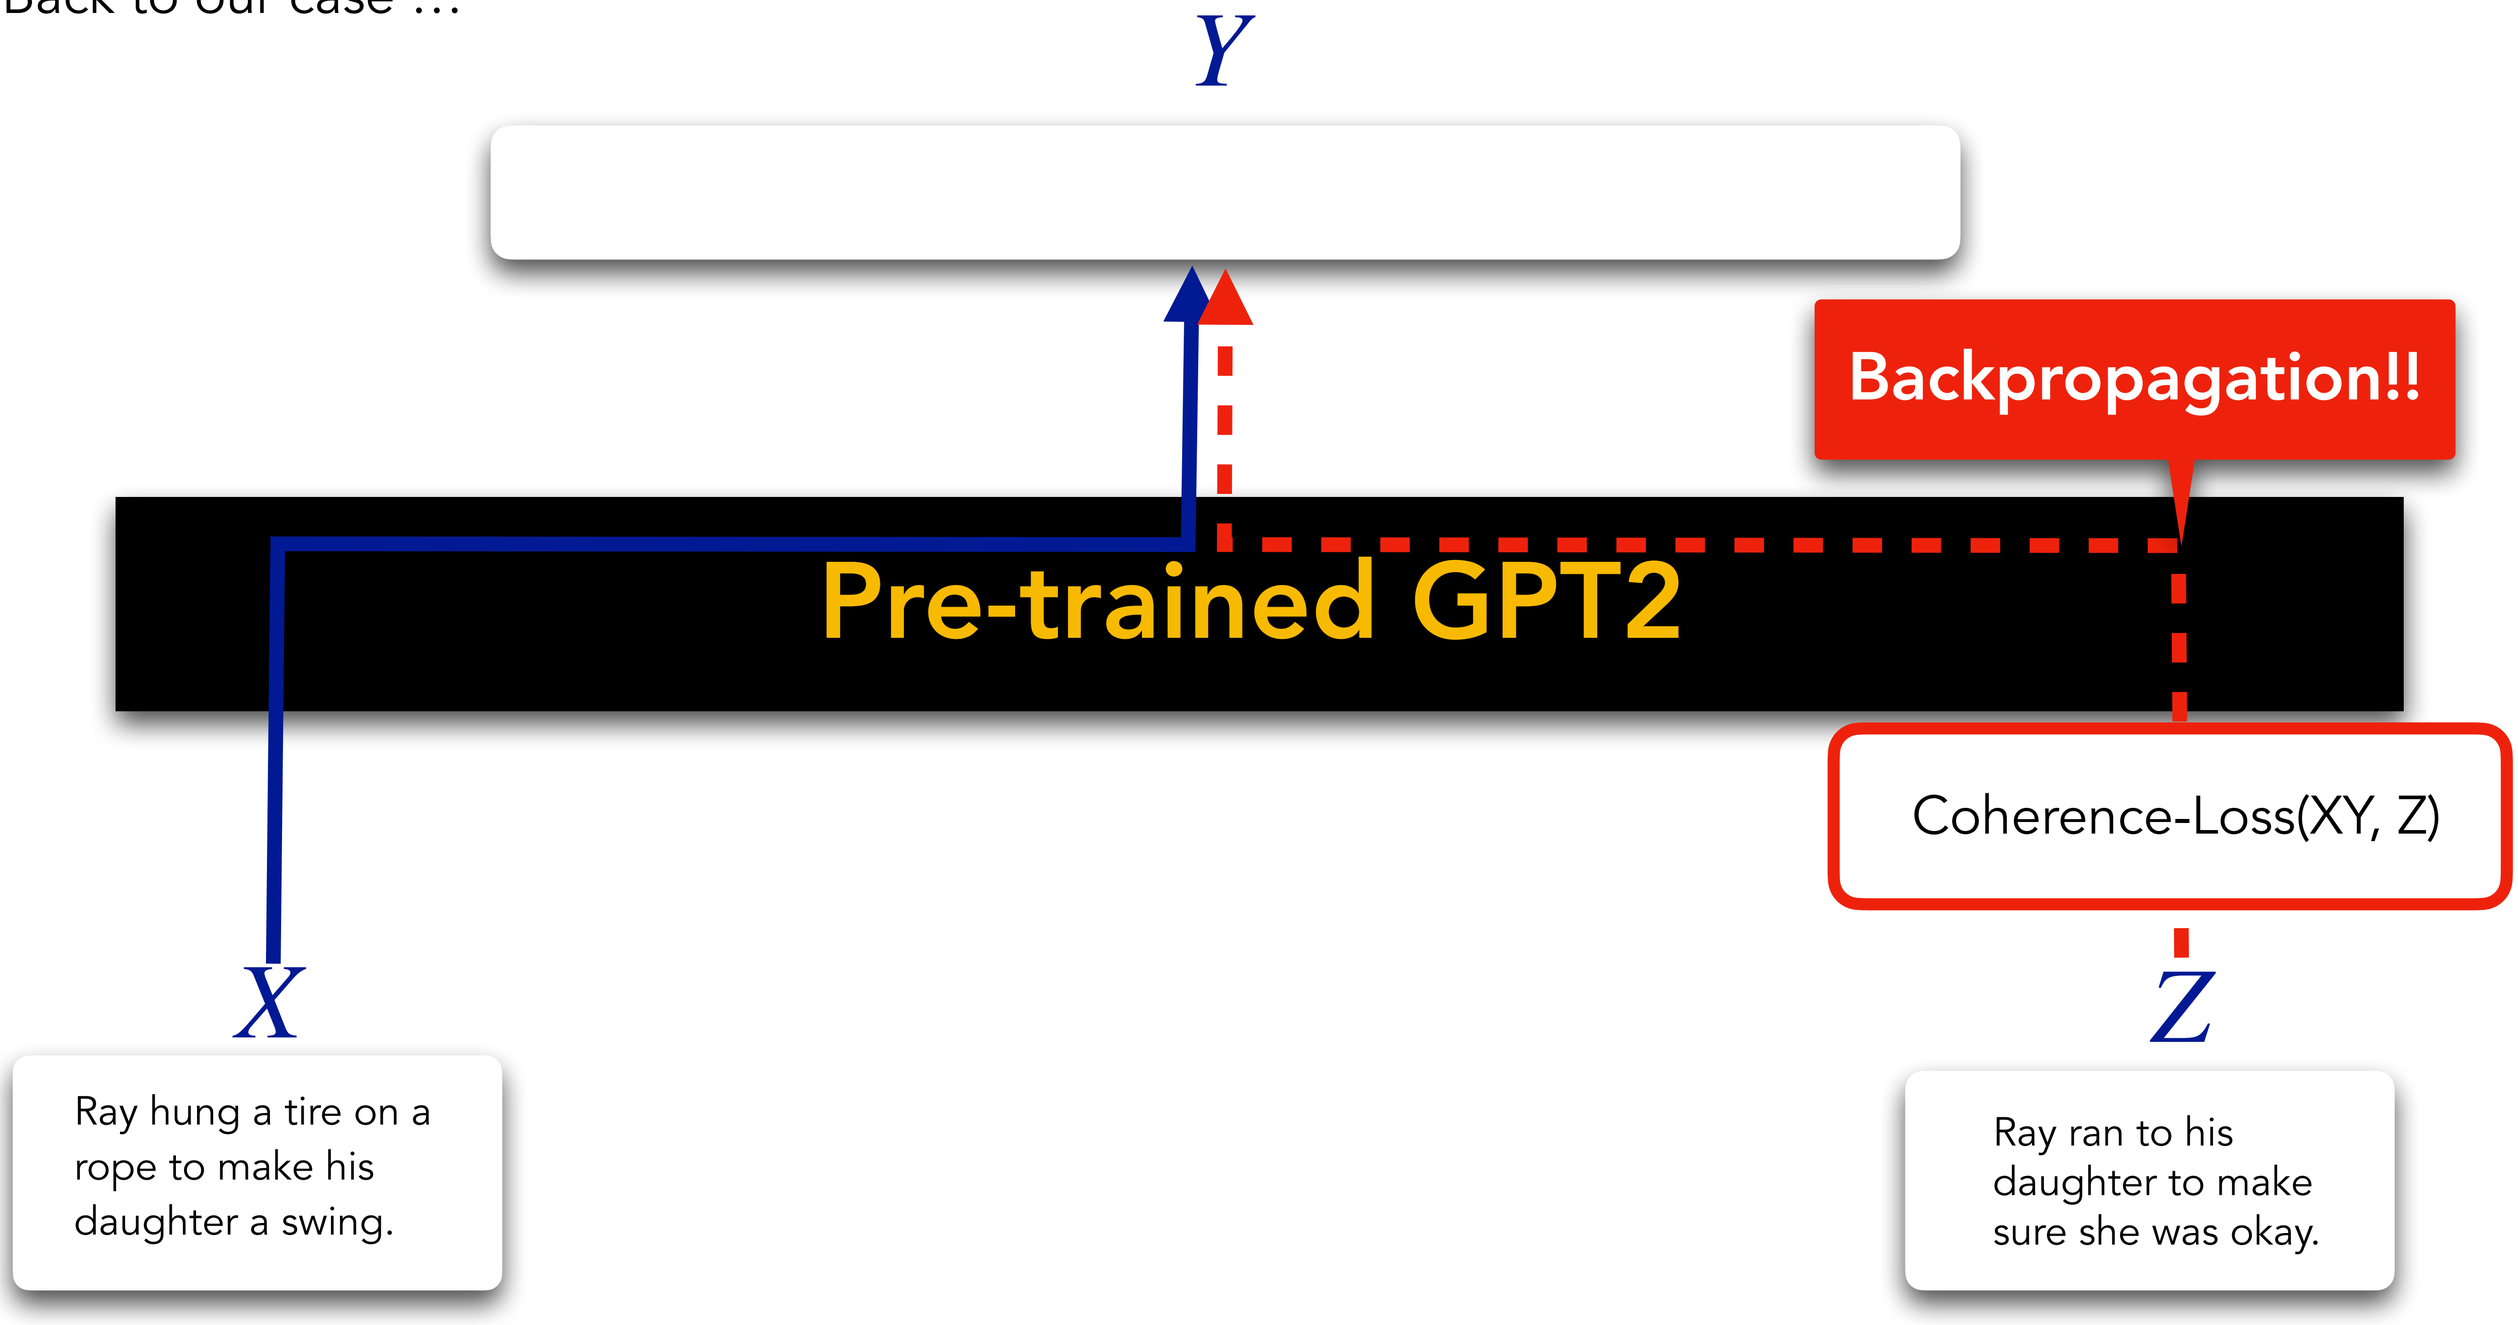

DELOREAN

(DEcoding for nonmonotonic LOgical REAsoNing)

Initialization

Just as how you do regular decoding

Y

$\tilde{y}_1$ $\tilde{y}_2$ ... $\tilde{y}_N$

LM

x_1 x_2 ... x_{N_X}

Ray hung a tire on a rope to make his daughter a swing.

Ray ran to his daughter to make sure she was okay.

X

Z

Backward Pass

Backpropagate future information

Y

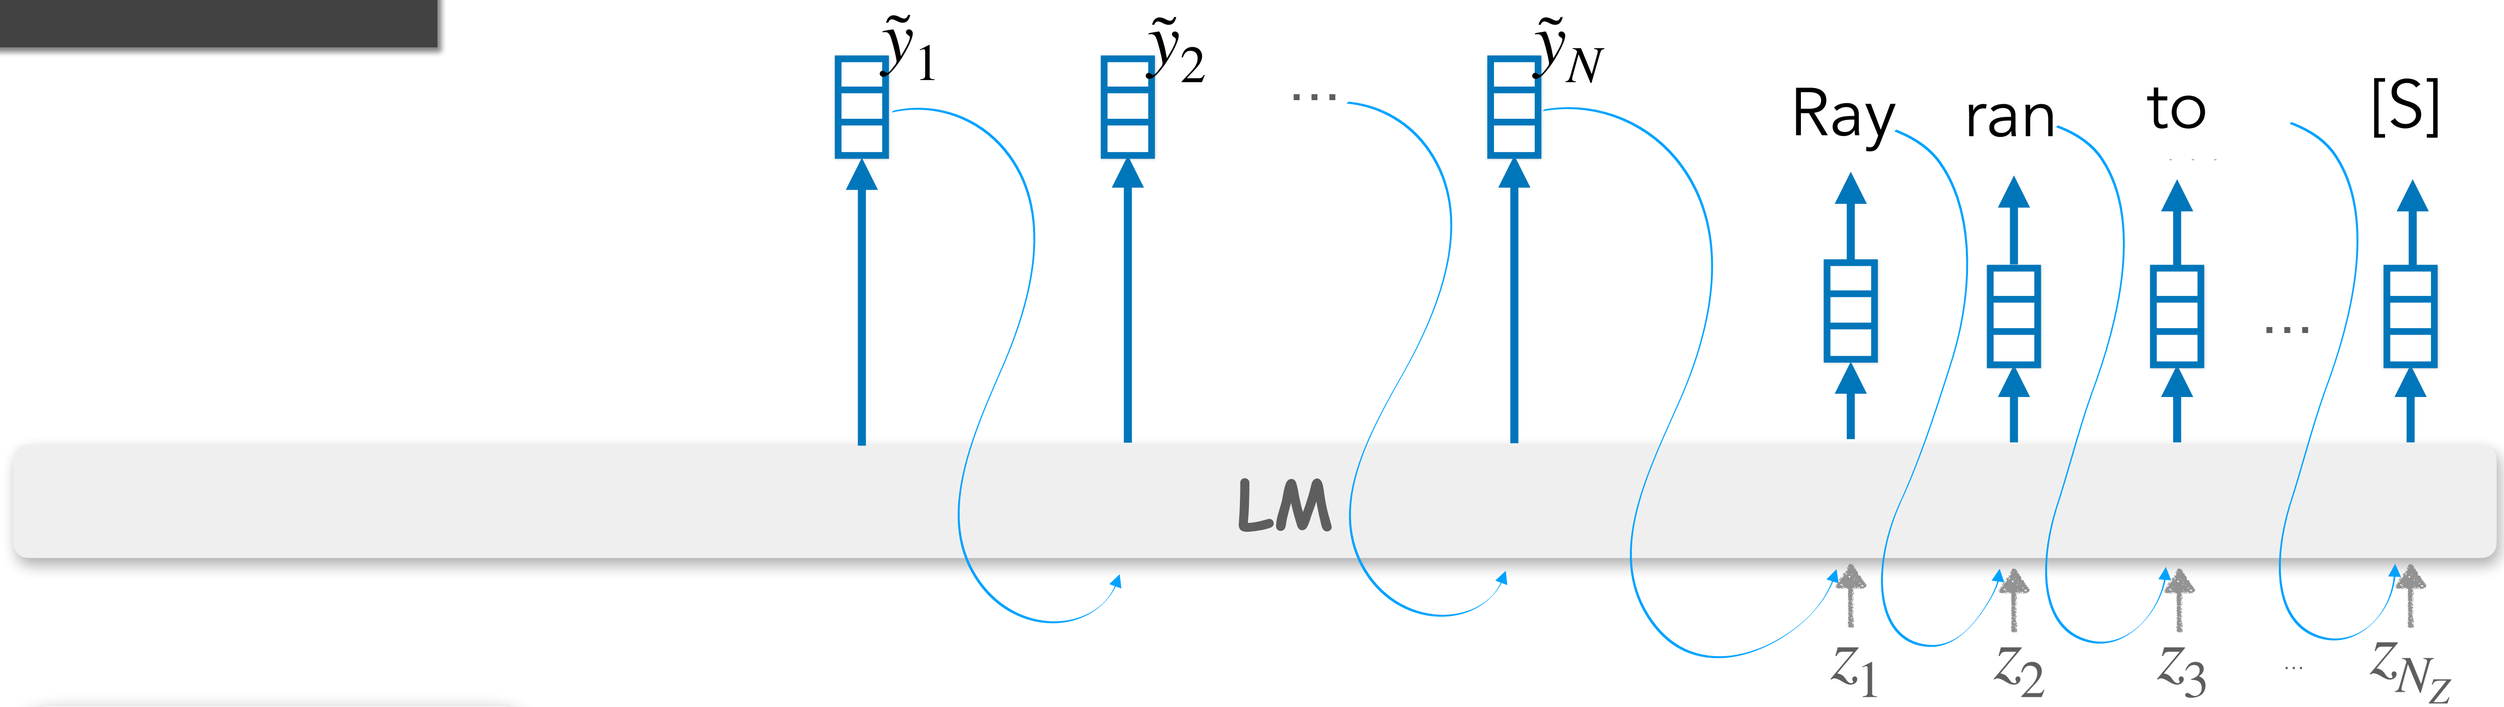

Ray hung a tire on a rope to make his daughter a swing.

Ray ran to his daughter to make sure she was okay.

X

Z

Backward Pass

Backpropagate future information

Loss(Z | X, Y)

Y

$\tilde{y}^b_1$ $\tilde{y}^b_2$ ... $\tilde{y}^b_N$

Backpropagation

$\tilde{y}_1$ $\tilde{y}_2$ ... $\tilde{y}_N$

$\mathcal{L}(X, \tilde{Y}, Z) := -\sum_{n=1}^{N_Z} \log P_{\text{LM}}(z_n | X, \tilde{Y}, Z_{1:n-1})$

Ray ran to ... [S]

LM

z_1 z_2 z_3 ... z_{N_Z}

Ray hung a tire on a rope to make his daughter a swing.

X

Ray ran to his daughter to make sure she was okay.

Z

Forward Pass

Mix both past and future information

Y

$\tilde{y}_1^b$ $\tilde{y}_2^b$ ... $\tilde{y}_N^b$

$\tilde{y}_1$ $\tilde{y}_2$ ... $\tilde{y}_N$

Backpropagation

$\mathcal{L}(X, \tilde{Y}, Z) := -\sum_{n=1}^{N_Z} \log P_{\mathrm{LM}}(z_n | X, \tilde{Y}, Z_{1:n-1})$

Ray ran to ... [S]

LM

z_1 z_2 z_3 ... z_{N_Z}

Ray hung a tire on a rope to make his daughter a swing.

X

Ray ran to his daughter to make sure she was okay.

Z

Forward Pass

Mix both past and future information

Y

$\tilde{y}_1^b$ $\tilde{y}_2^b$ ... $\tilde{y}_N^b$

$\tilde{y}_1$ $\tilde{y}_2$ ... $\tilde{y}_N$

$\tilde{y}_1^f$ $\tilde{y}_2^f$ ... $\tilde{y}_N^f$

Backpropagation

$\mathcal{L}(X, \tilde{Y}, Z) := -\sum_{n=1}^{N_Z} \log P_{\mathrm{LM}}(z_n | X, \tilde{Y}, Z_{1:n-1})$

Ray ran to ... [S]

LM

x_1 x_2 ... x_{N_X}

z_1 z_2 z_3 ... z_{N_Z}

Ray hung a tire on a rope to make his daughter a swing.

Ray ran to his daughter to make sure she was okay.

X

Z

Repeat *T* times

Y

Backpropagation

$\mathcal{L}(X, \tilde{Y}, Z) := -\sum_{n=1}^{N_Z} \log P_{\mathrm{LM}}(z_n | X, \tilde{Y}, Z_{1:n-1})$

$\tilde{y}_1^b$ $\tilde{y}_2^b$ … $\tilde{y}_N^b$

$\tilde{y}_1$ $\tilde{y}_2$ … $\tilde{y}_N$

$\tilde{y}_1^f$ $\tilde{y}_2^f$ … $\tilde{y}_N^f$

Ray ran to … [S]

LM

x_1 x_2 … x_{N_X}

z_1 z_2 z_3 … z_{N_Z}

Ray hung a tire on a rope to make his daughter a swing.

Ray ran to his daughter to make sure she was okay.

X

Z

Sampling

Y

Output: She hit the rope and the tire fell on top of her.

Backpropagation

$\mathcal{L}(X, \tilde{Y}, Z) := -\sum_{n=1}^{N_Z} \log P_{\mathrm{LM}}(z_n | X, \tilde{Y}, Z_{1:n-1})$

$\tilde{y}_1^b$ $\tilde{y}_2^b$ … $\tilde{y}_N^b$

$\tilde{y}_1$ $\tilde{y}_2$ … $\tilde{y}_N$

$\tilde{y}_1^f$ $\tilde{y}_2^f$ … $\tilde{y}_N^f$

Ray ran to … [S]

LM

x_1 x_2 … x_{N_X}

z_1 z_2 z_3 … z_{N_Z}

Ray hung a tire on a rope to make his daughter a swing.

Ray ran to his daughter to make sure she was okay.

X

Z

Counterfactual Reasoning?

Y

Backpropagation

?

$\tilde{y}_1^b$ $\tilde{y}_2^b$ … $\tilde{y}_N^b$

$\tilde{y}_1$ $\tilde{y}_2$ … $\tilde{y}_N$

$\tilde{y}_1^f$ $\tilde{y}_2^f$ … $\tilde{y}_N^f$

LM

x_1 x_2 … x_{N_X}

z_1 z_2 z_3 … z_{N_Z}

X

Z

Counterfactual Reasoning?

Distance-Loss: (Y, Z)

Y

Backpropagation

$\mathcal{L}(X, \tilde{Y}, Z) := \mathrm{KL}(Z \| \mathrm{softmax}(\tilde{Y}/\tau))$

?

$\tilde{y}_1^b$ $\tilde{y}_2^b$ ... $\tilde{y}_N^b$

$\tilde{y}_1$ $\tilde{y}_2$ ... $\tilde{y}_N$

$\tilde{y}_1^f$ $\tilde{y}_2^f$ ... $\tilde{y}_N^f$

LM

x_1 x_2 ... x_{N_X}

z_1 z_2 z_3 ... z_{N_Z}

X

Z

Counterfactual Reasoning?

Distance-Loss: (Y, Z)

Y

$\tilde{y}^b_1$ $\tilde{y}^b_2$ … $\tilde{y}^b_N$

$\tilde{y}_1$ $\tilde{y}_2$ … $\tilde{y}_N$

$\tilde{y}^f_1$ $\tilde{y}^f_2$ … $\tilde{y}^f_N$

Backpropagation

$$\mathcal{L}(X, \tilde{Y}, Z) := \mathrm{KL}\left(Z \| \mathrm{softmax}(\tilde{Y}/\tau)\right)$$

Zeke thought about … [S]

LM

x_1 x_2 … x_{N_X}

z_1 z_2 z_3 … z_{N_Z}

Zeke was throwing a party. [Counterfactual] All his friends were dressing up for this Game of Thrones themed party.

X

Zeke thought about being a vampire or a wizard. Then he decided on a scarier costume. Zeke dressed up like a skeleton.

Z

Counterfactual Reasoning?

Distance-Loss: (Y, Z)

Y

Zeke thought about Lannister, but he didn't want to look like a Lannister. He wanted to look like a Stark.Zeke dressed up like a *Stark*.

Backpropagation

$$\mathcal{L}(X, \tilde{Y}, Z) := \mathrm{KL}\left(Z \| \mathrm{softmax}(\tilde{Y}/\tau)\right)$$

$\tilde{y}^b_1$ $\tilde{y}^b_2$ … $\tilde{y}^b_N$

$\tilde{y}_1$ $\tilde{y}_2$ … $\tilde{y}_N$

$\tilde{y}^f_1$ $\tilde{y}^f_2$ … $\tilde{y}^f_N$

Zeke thought about … [S]

LM

x_1 x_2 … x_{N_X}

z_1 z_2 z_3 … z_{N_Z}

Zeke was throwing a party. [Counterfactual] All his friends were dressing up for this Game of Thrones themed party.

X

Zeke thought about being a vampire or a wizard. Then he decided on a scarier costume. Zeke dressed up like a skeleton.

Z

Abductive Reasoning

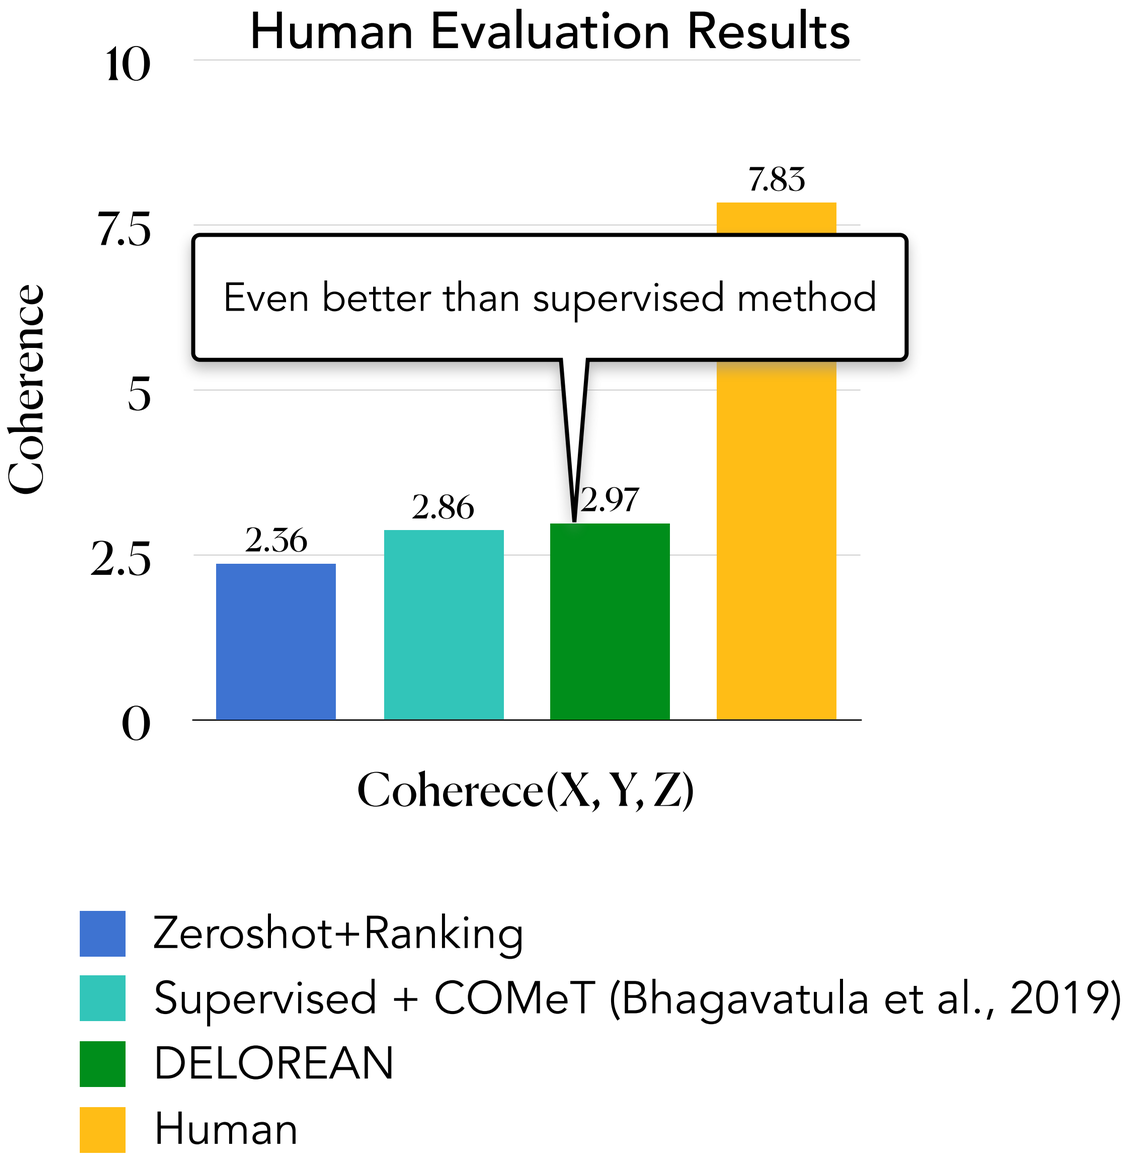

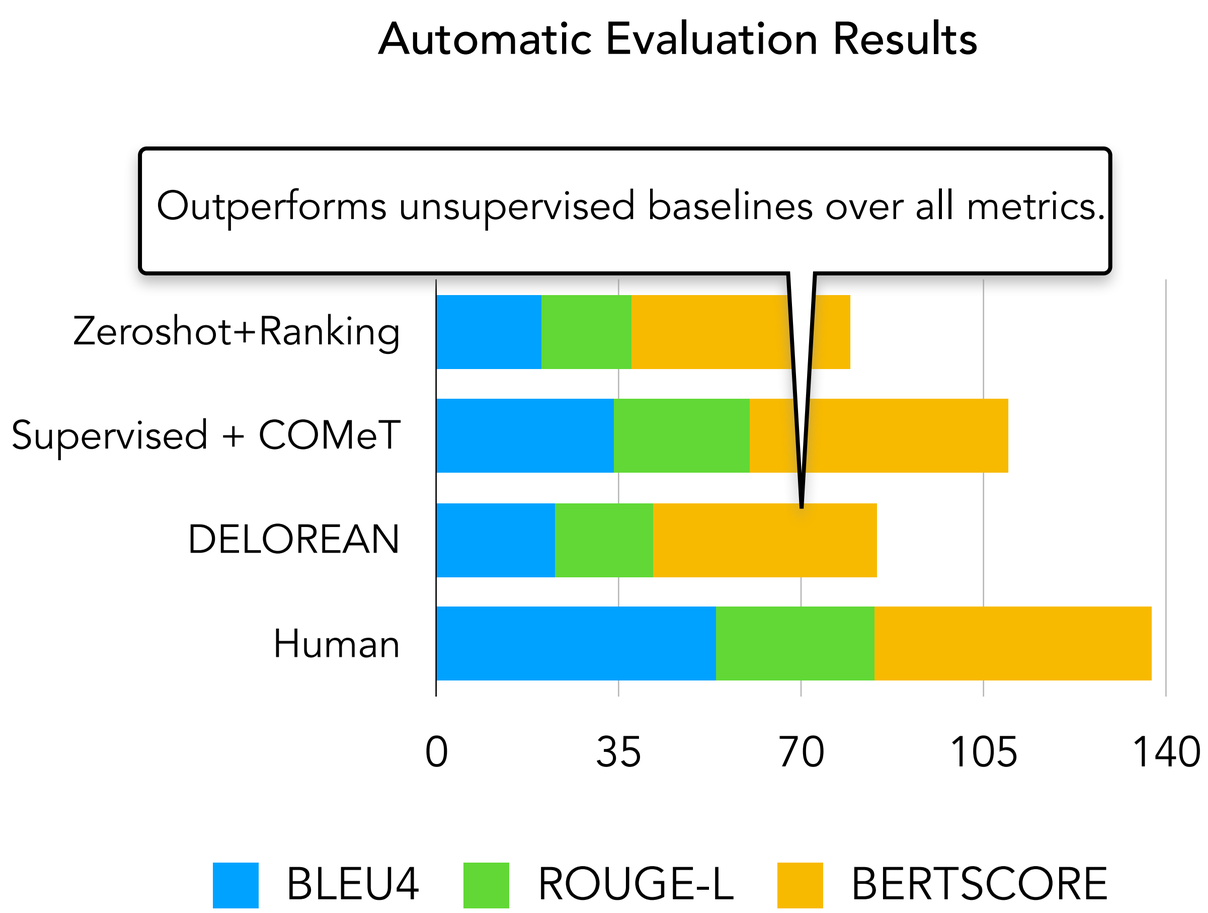

Please check the paper for more baselines …

Language

Induction | Deduction | Abduction | Counterfactual

Reasoning

Reasoning as ***generative tasks***

- As opposed to ***discriminative tasks*** (i.e., categorization)
- Because the space of reasoning in language is ***infinite***

"thinking out loud"
We often think as we speak, on the fly, word-by-word without enumerating all possible alternative sentences

Reasoning serves the purpose of **communication**

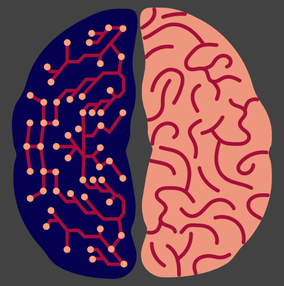

NeuroLogic Decoding

(Un)supervised Neural Text Generation with **Predicate Logic Constraints**

NAACL 2021

Ximing Lu

Peter West

Rowan Zellers

Ronan LeBras

Chandra Bhagavatula

Yejin Choi

Seq2Seq

Machine Translation

X

The physician told the baker that she had cancer.

Y

Der Arzt sagte dem Bäckerin, dass er Krebs habe.

Dialogue Response

X

type	hotel
count	182
dogs allowed	don't care

Y

There are 182 hotels if you do not care whether dogs are allowed .

Y

Language Model

X

COMMONGEN

X

{ food, table, sit, front }

Y

The man sat with his food at the front of the table.

Image Captioning

X

Y

Man in blue wetsuit is surfing on wave.

CommonGen (Lin et al. EMNLP 2020)

missing keyword { lose, ride }

A man is trying to keep his balance as he falls off a board .

should use **all** given keywords ....

Fine-tuned Language Model

X

{ board, lose, ride, fall, balance }

NEUROLOGIC DECODING

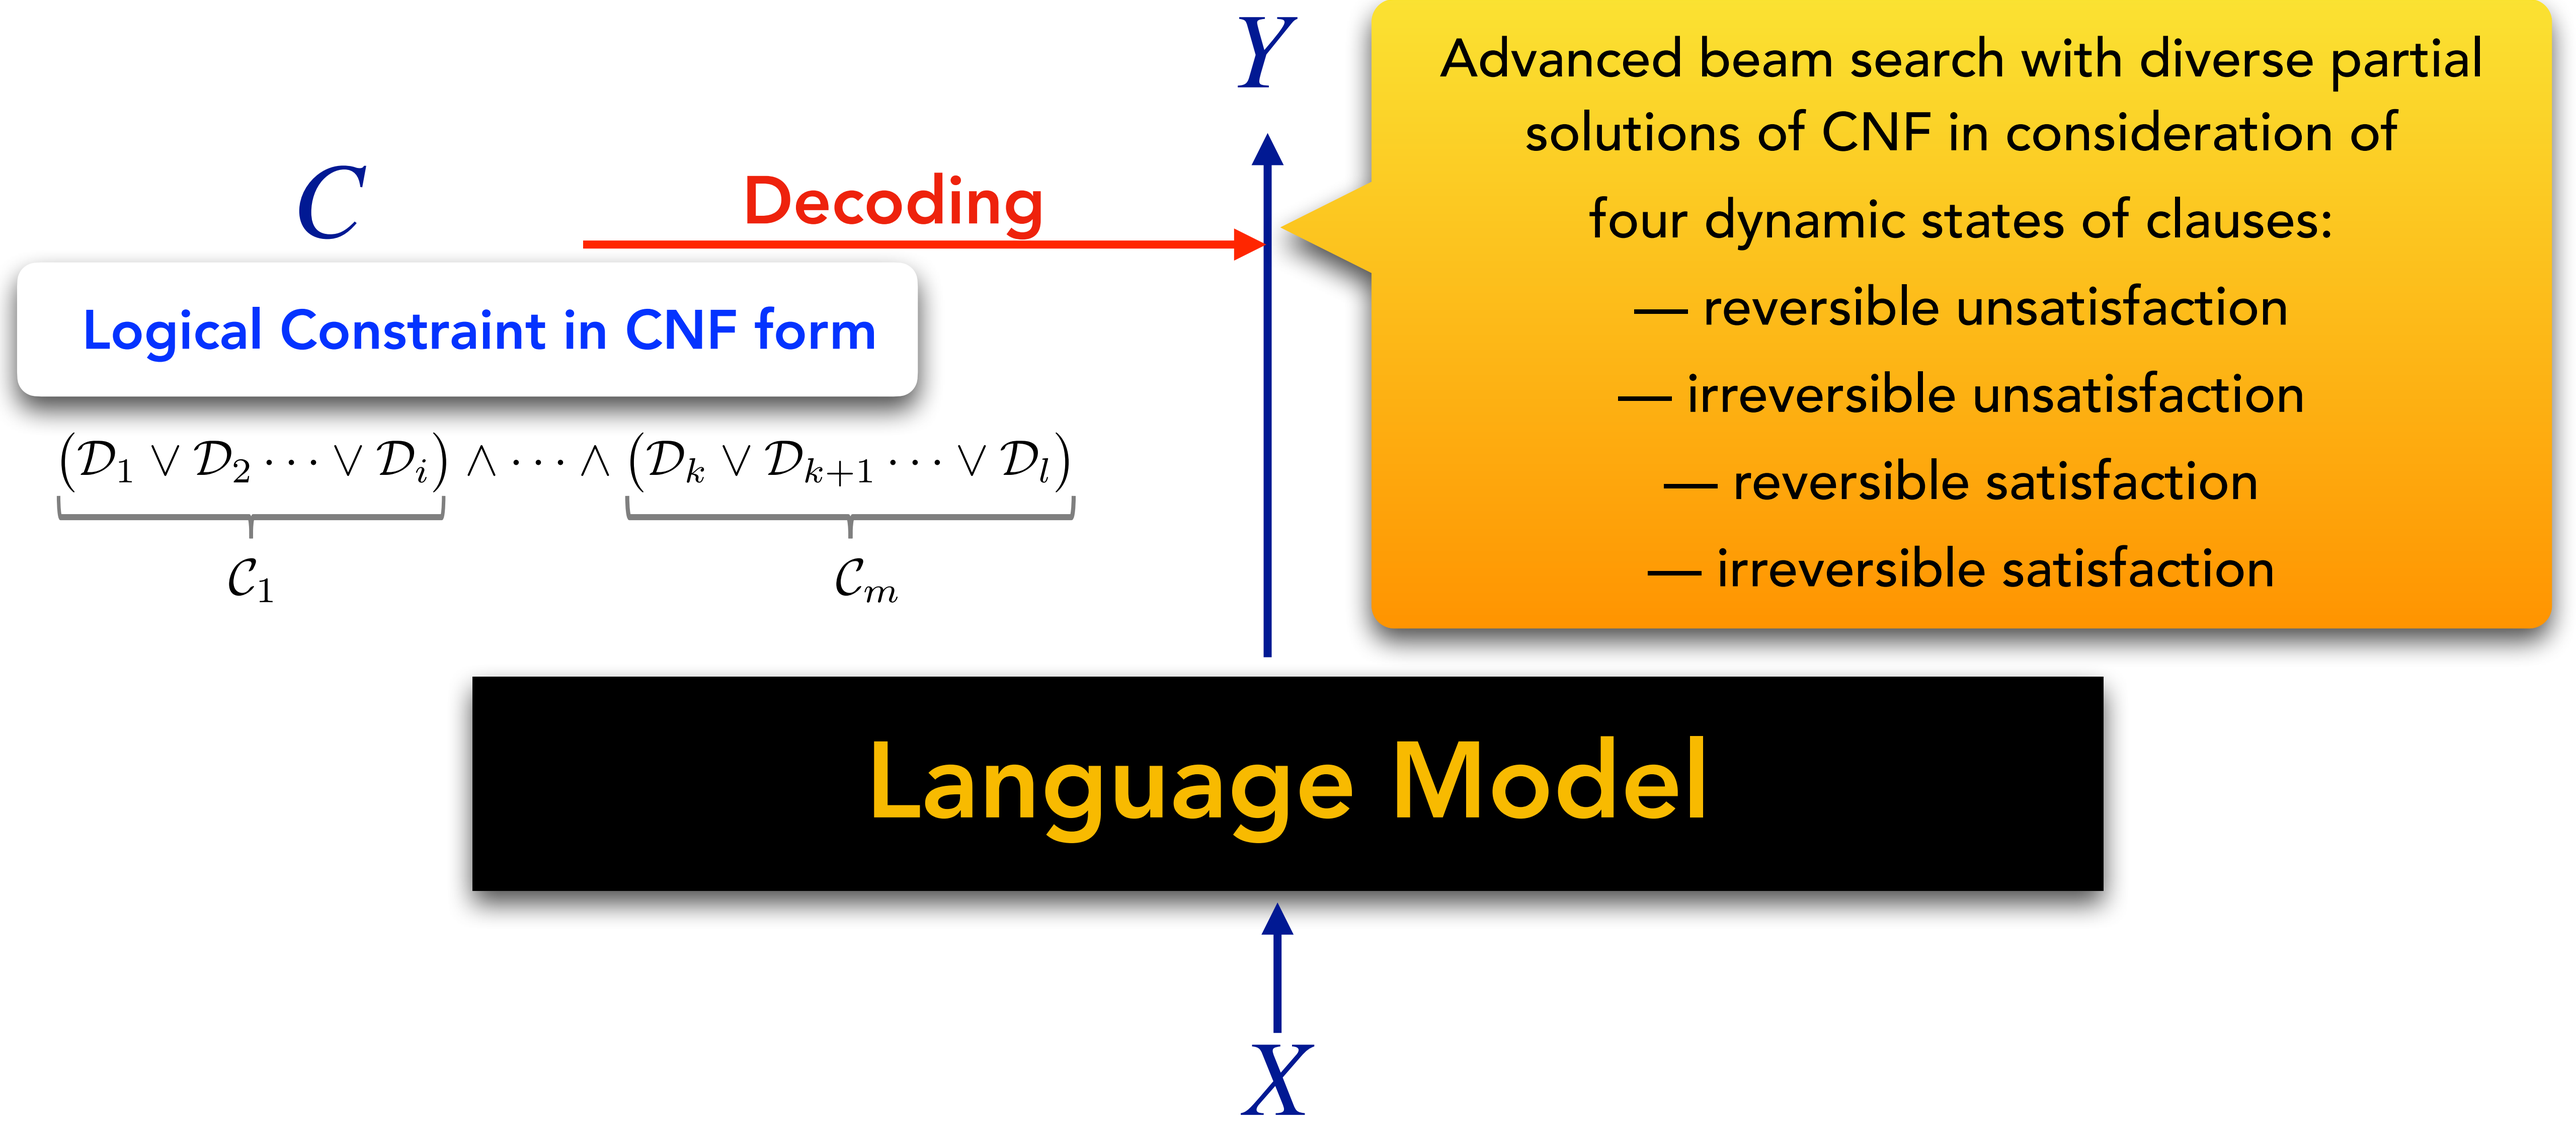

CommonGen (Zero-shot)

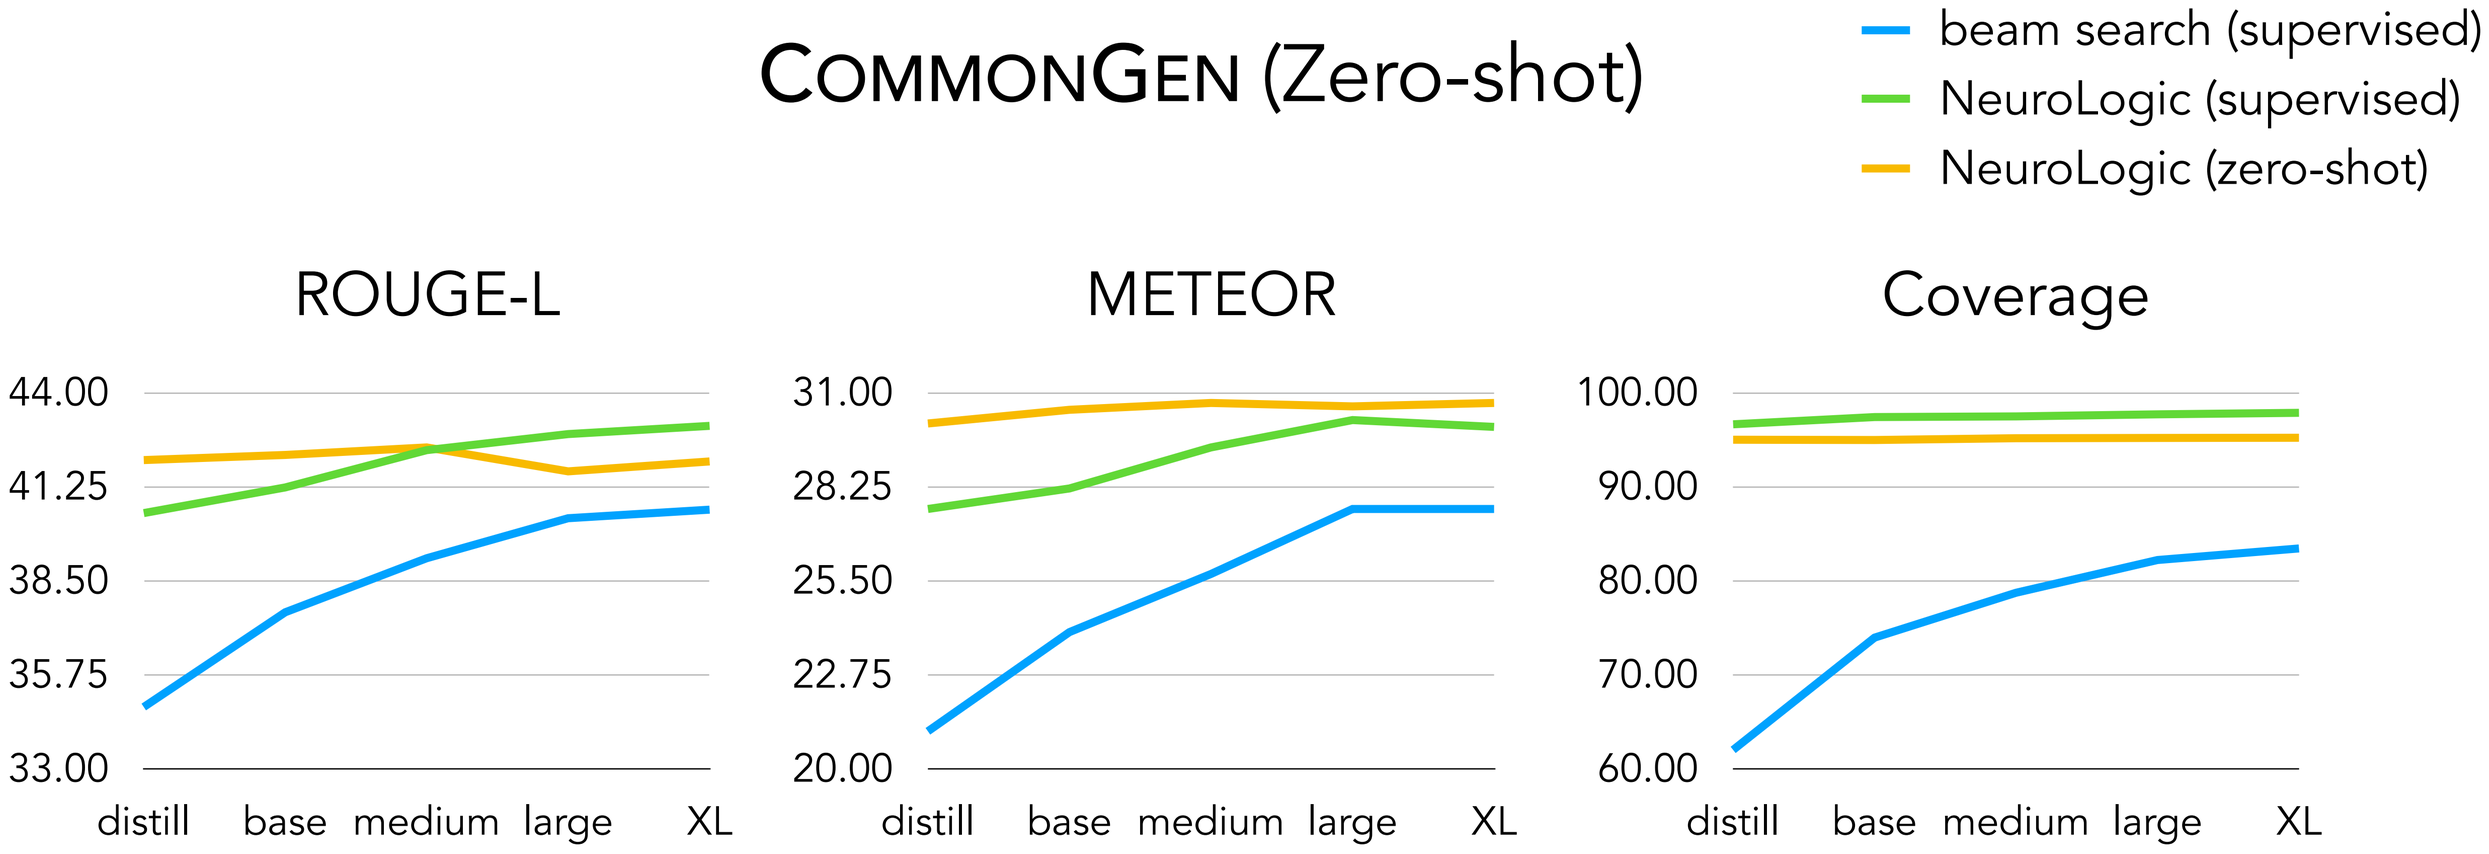

Unsupervised NeuroLogic outperforms **supervised approaches**

Unsupervised NeuroLogic on smaller networks outperforms **supervised approaches on larger networks!**

Path to commonsense?

Obligatory

🔥Controversial Remarks🔥 of the Day

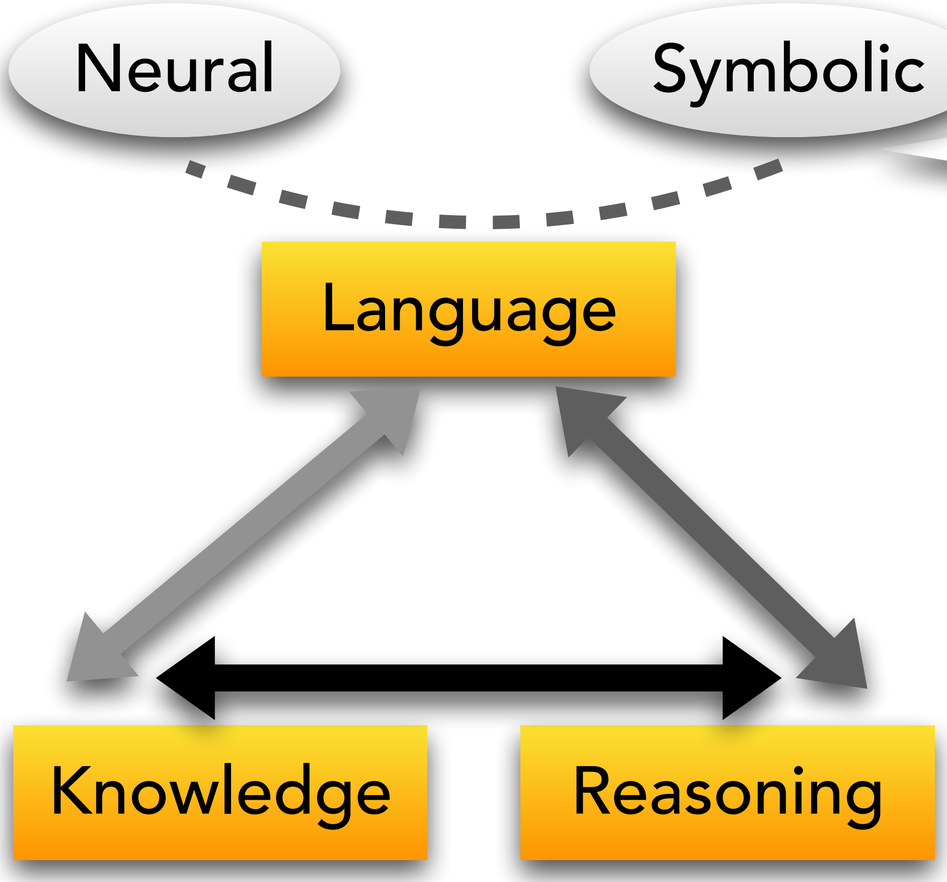

1. the continuum between **knowledge** and **reasoning**
2. the interplay between **reasoning** and **language generation**
3. the blend between **neural** vs **symbolic** representation

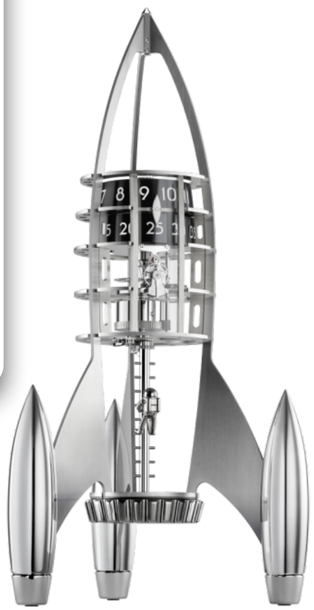

Path to commonsense?

Obligatory

🔥Controversial Remarks🔥 of the Day

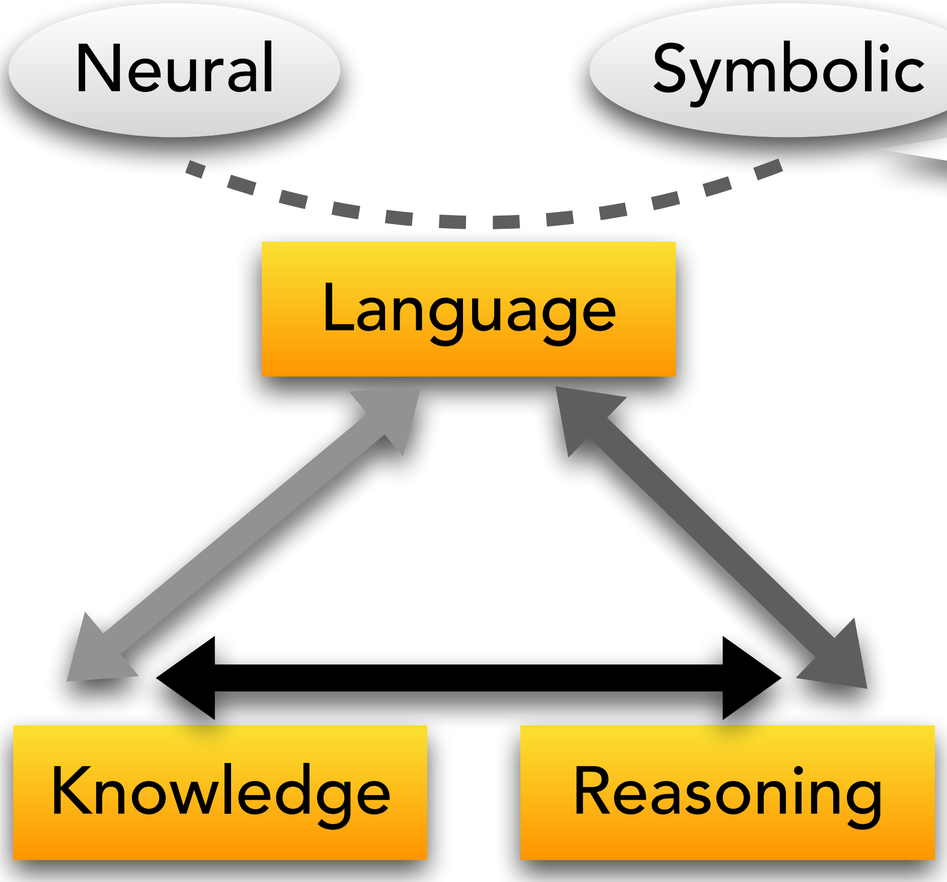

1. the continuum between **knowledge** and **reasoning**
2. the interplay between **reasoning** and **language generation**
3. the blend between **neural** vs **symbolic** representation

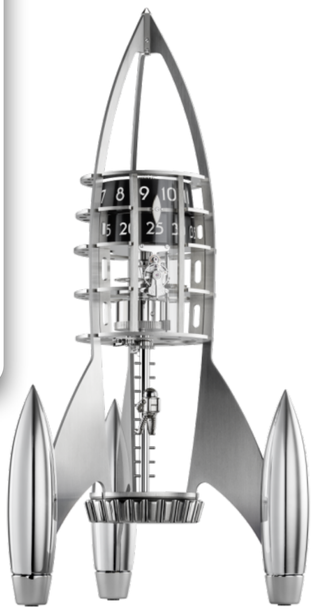

(COMET-) ATOMIC$^{20}_{20}$:

On Symbolic and Neural Commonsense Knowledge Graphs

— wait, doesn't GPT-3 know everything? —

AAAI 2021

Jena Hwang

Chandra Bhagavatula

Ronan Le Bras

Jeff Da

Keisuke Sakaguchi

Antoine Bosseult

Me

Language models != knowledge models

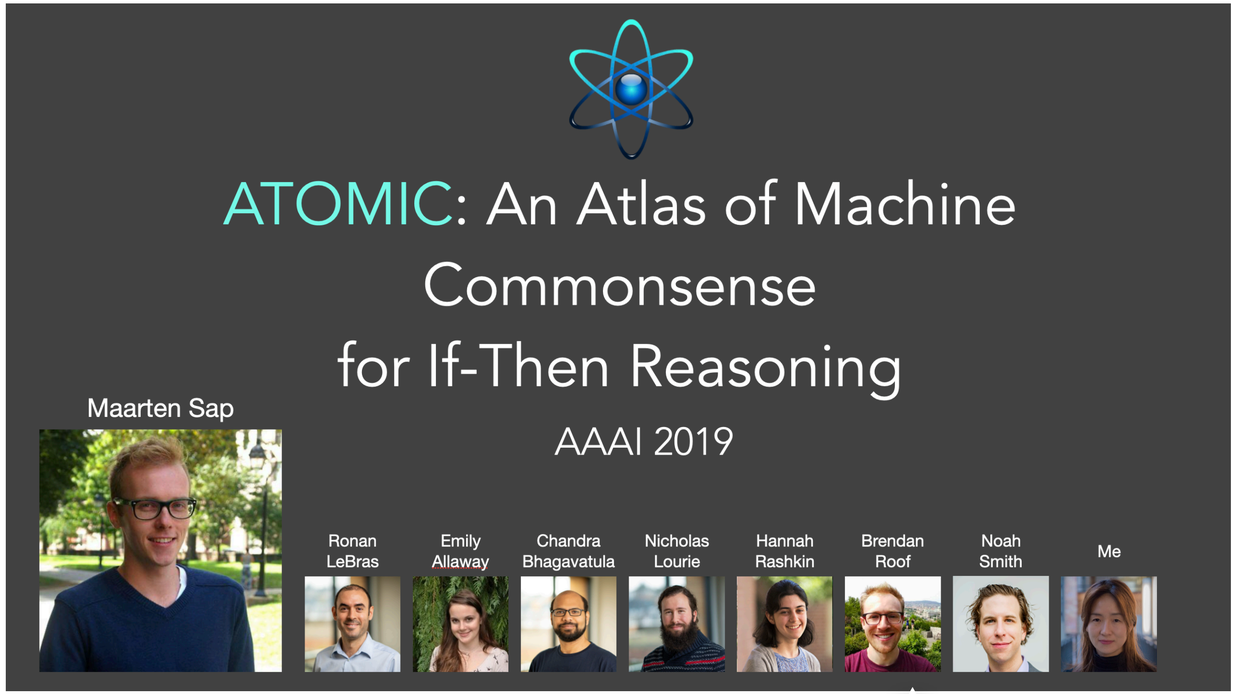

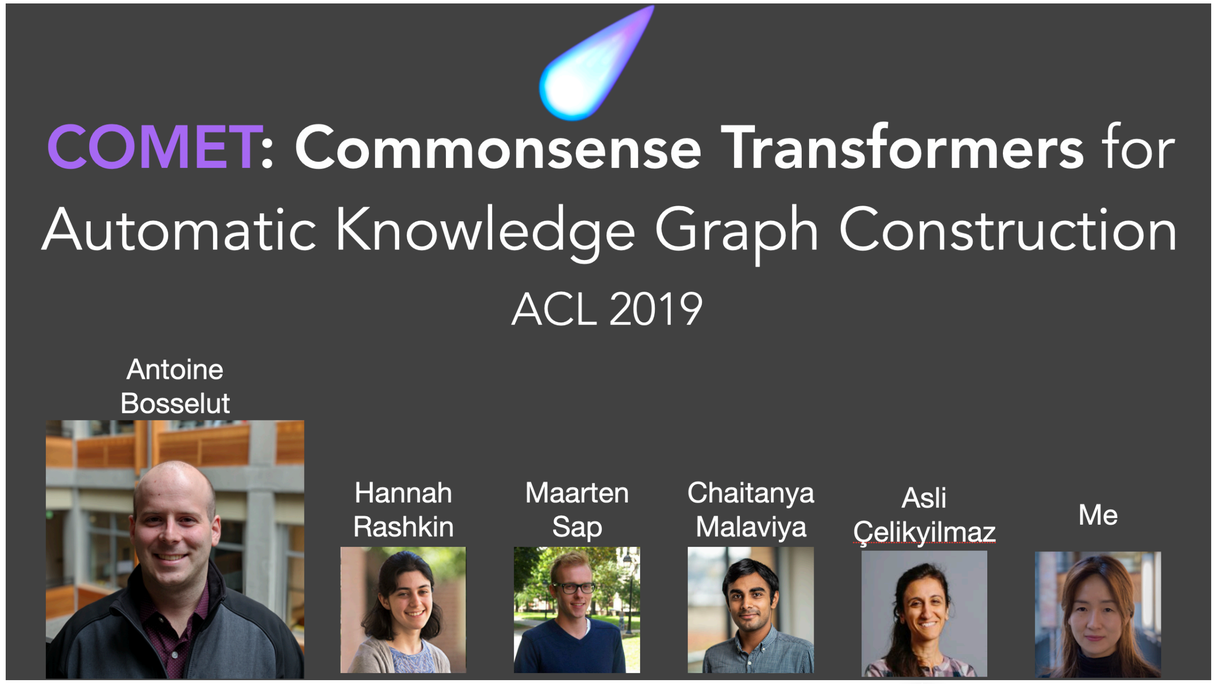

Symbolic commonsense knowledge graph

Neural commonsense model

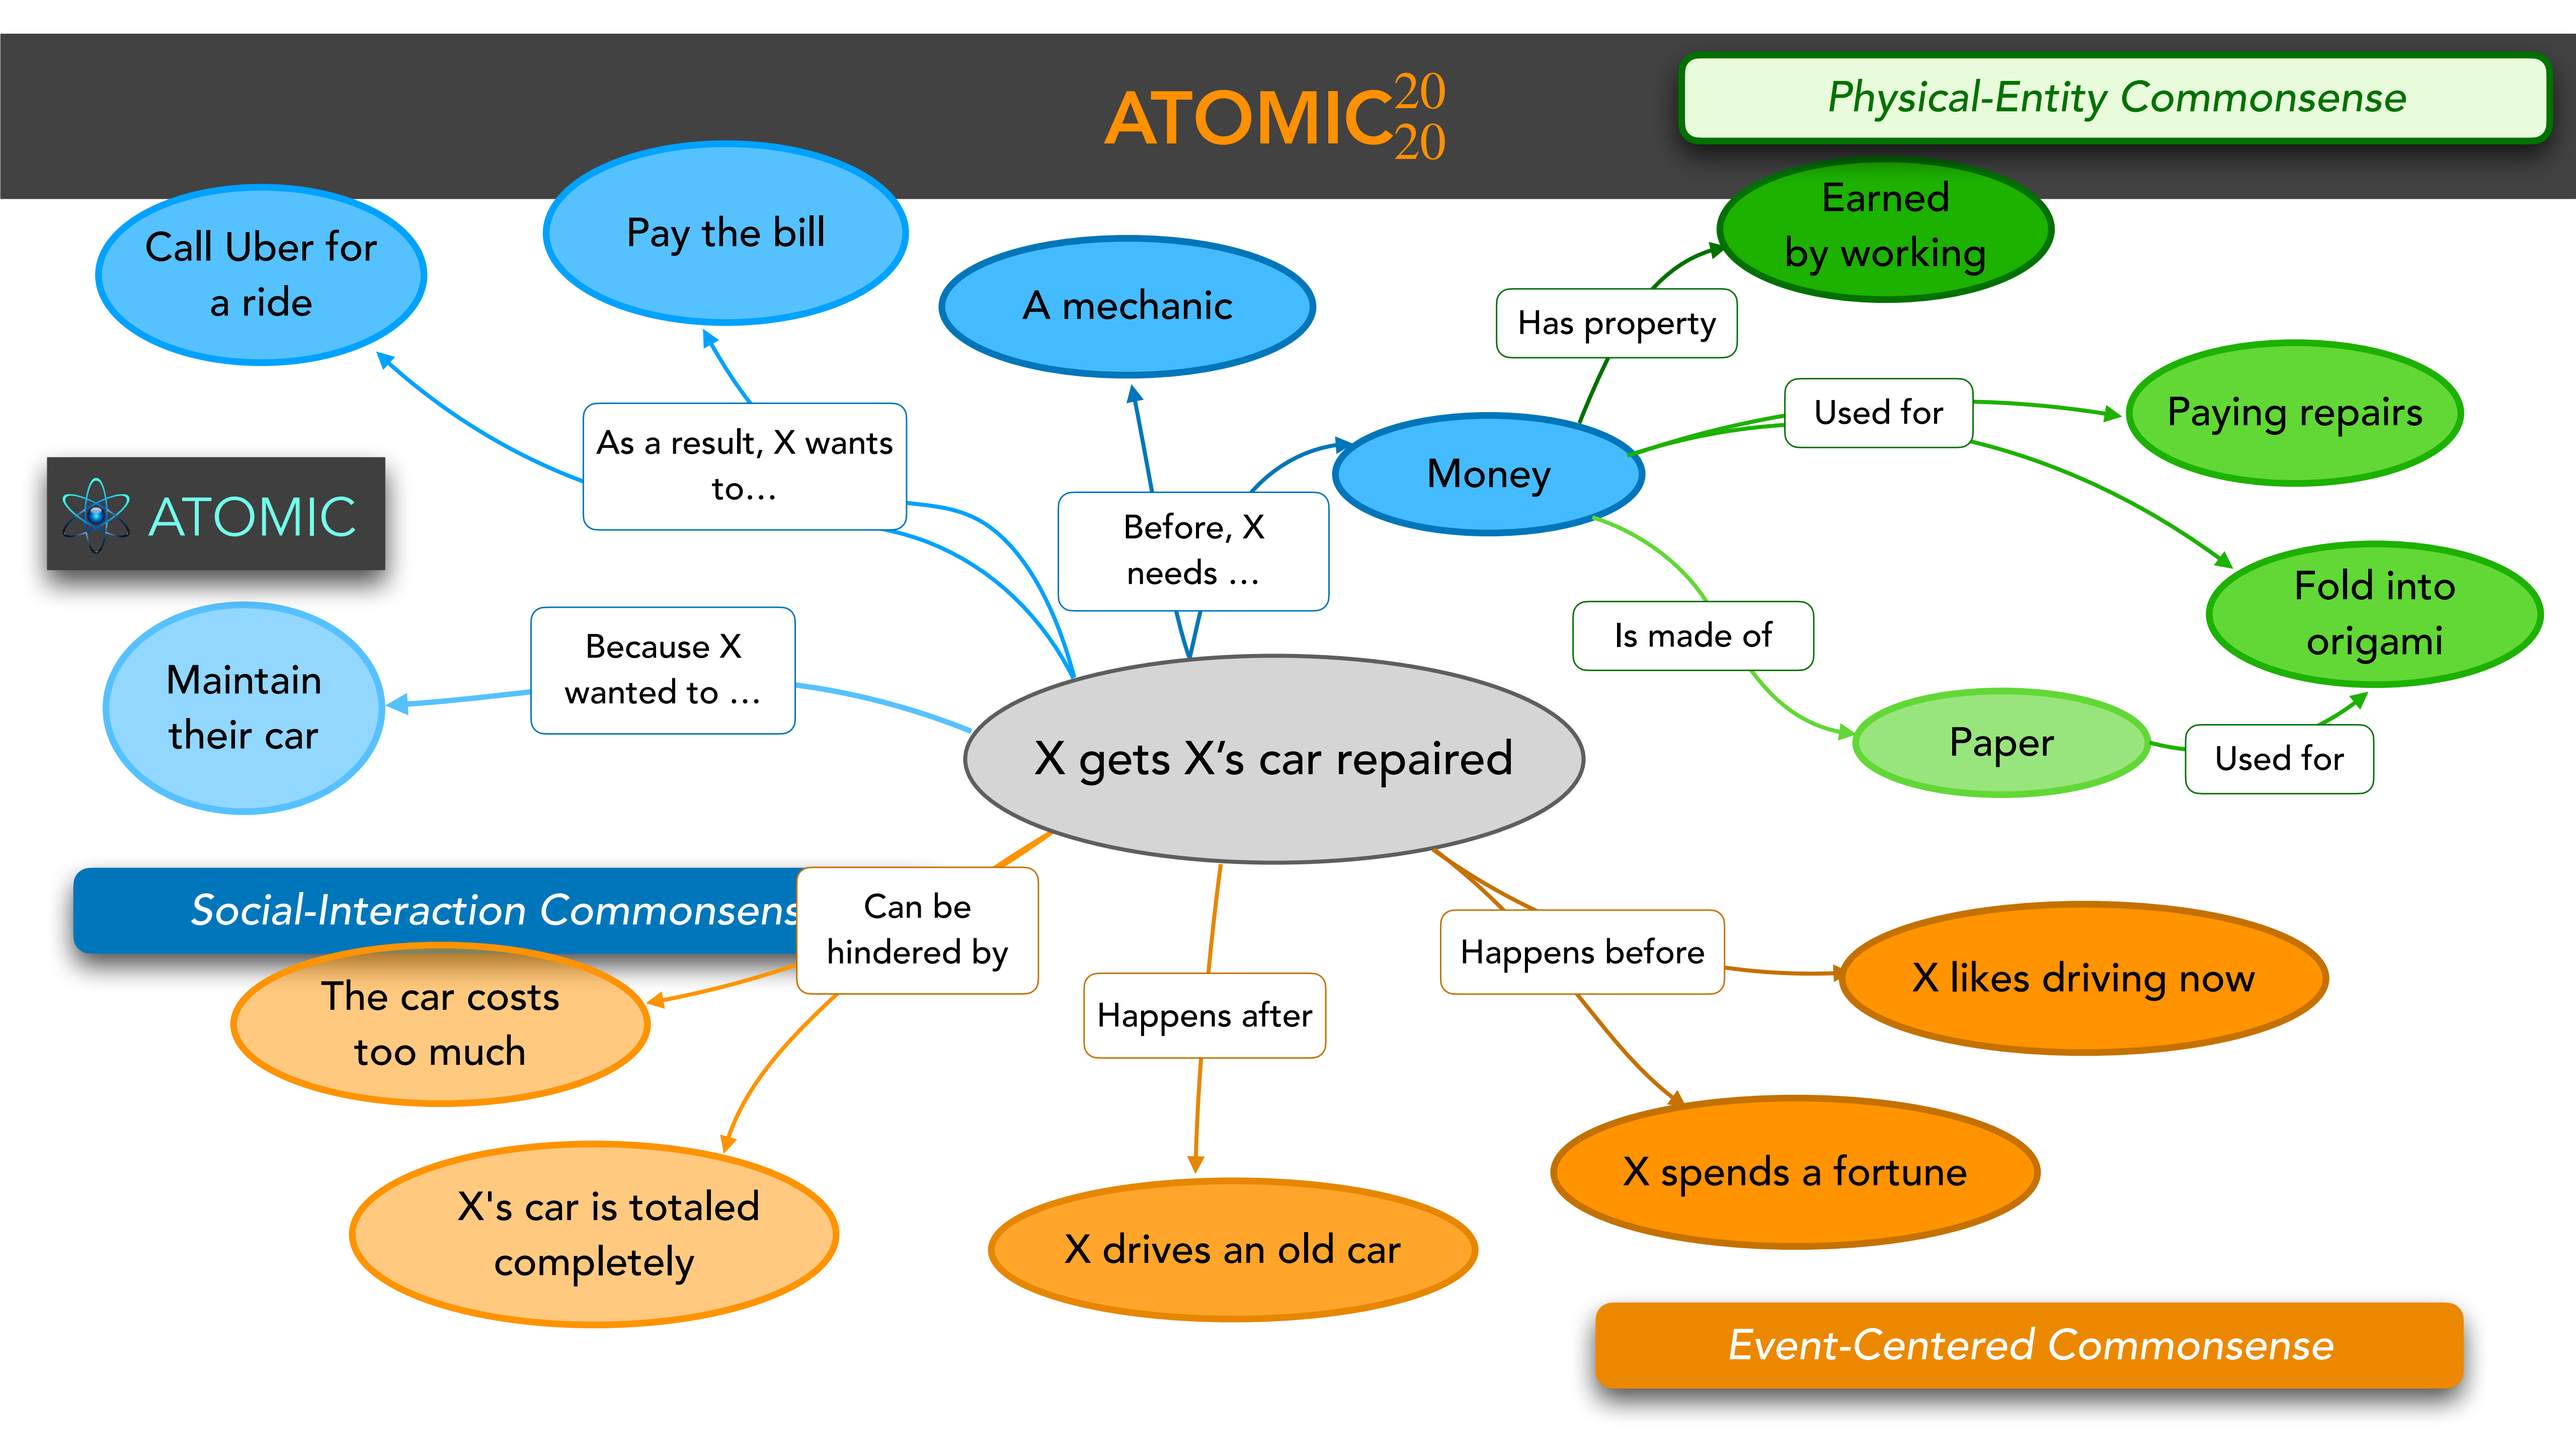
ATOMIC 2020
Physical-Entity Commonsense
Call Uber for a ride
Pay the bill
A mechanic
Earned by working
Has property
As a result, X wants to…
Before, X needs …
Money
Used for
Paying repairs
ATOMIC
Fold into origami
Maintain their car
Because X wanted to …
Is made of
X gets X's car repaired
Paper
Used for
Social-Interaction Commonsens
Can be hindered by
Happens before
X likes driving now
The car costs too much
Happens after
X spends a fortune
X's car is totaled completely
X drives an old car
Event-Centered Commonsense

ATOMIC 2020

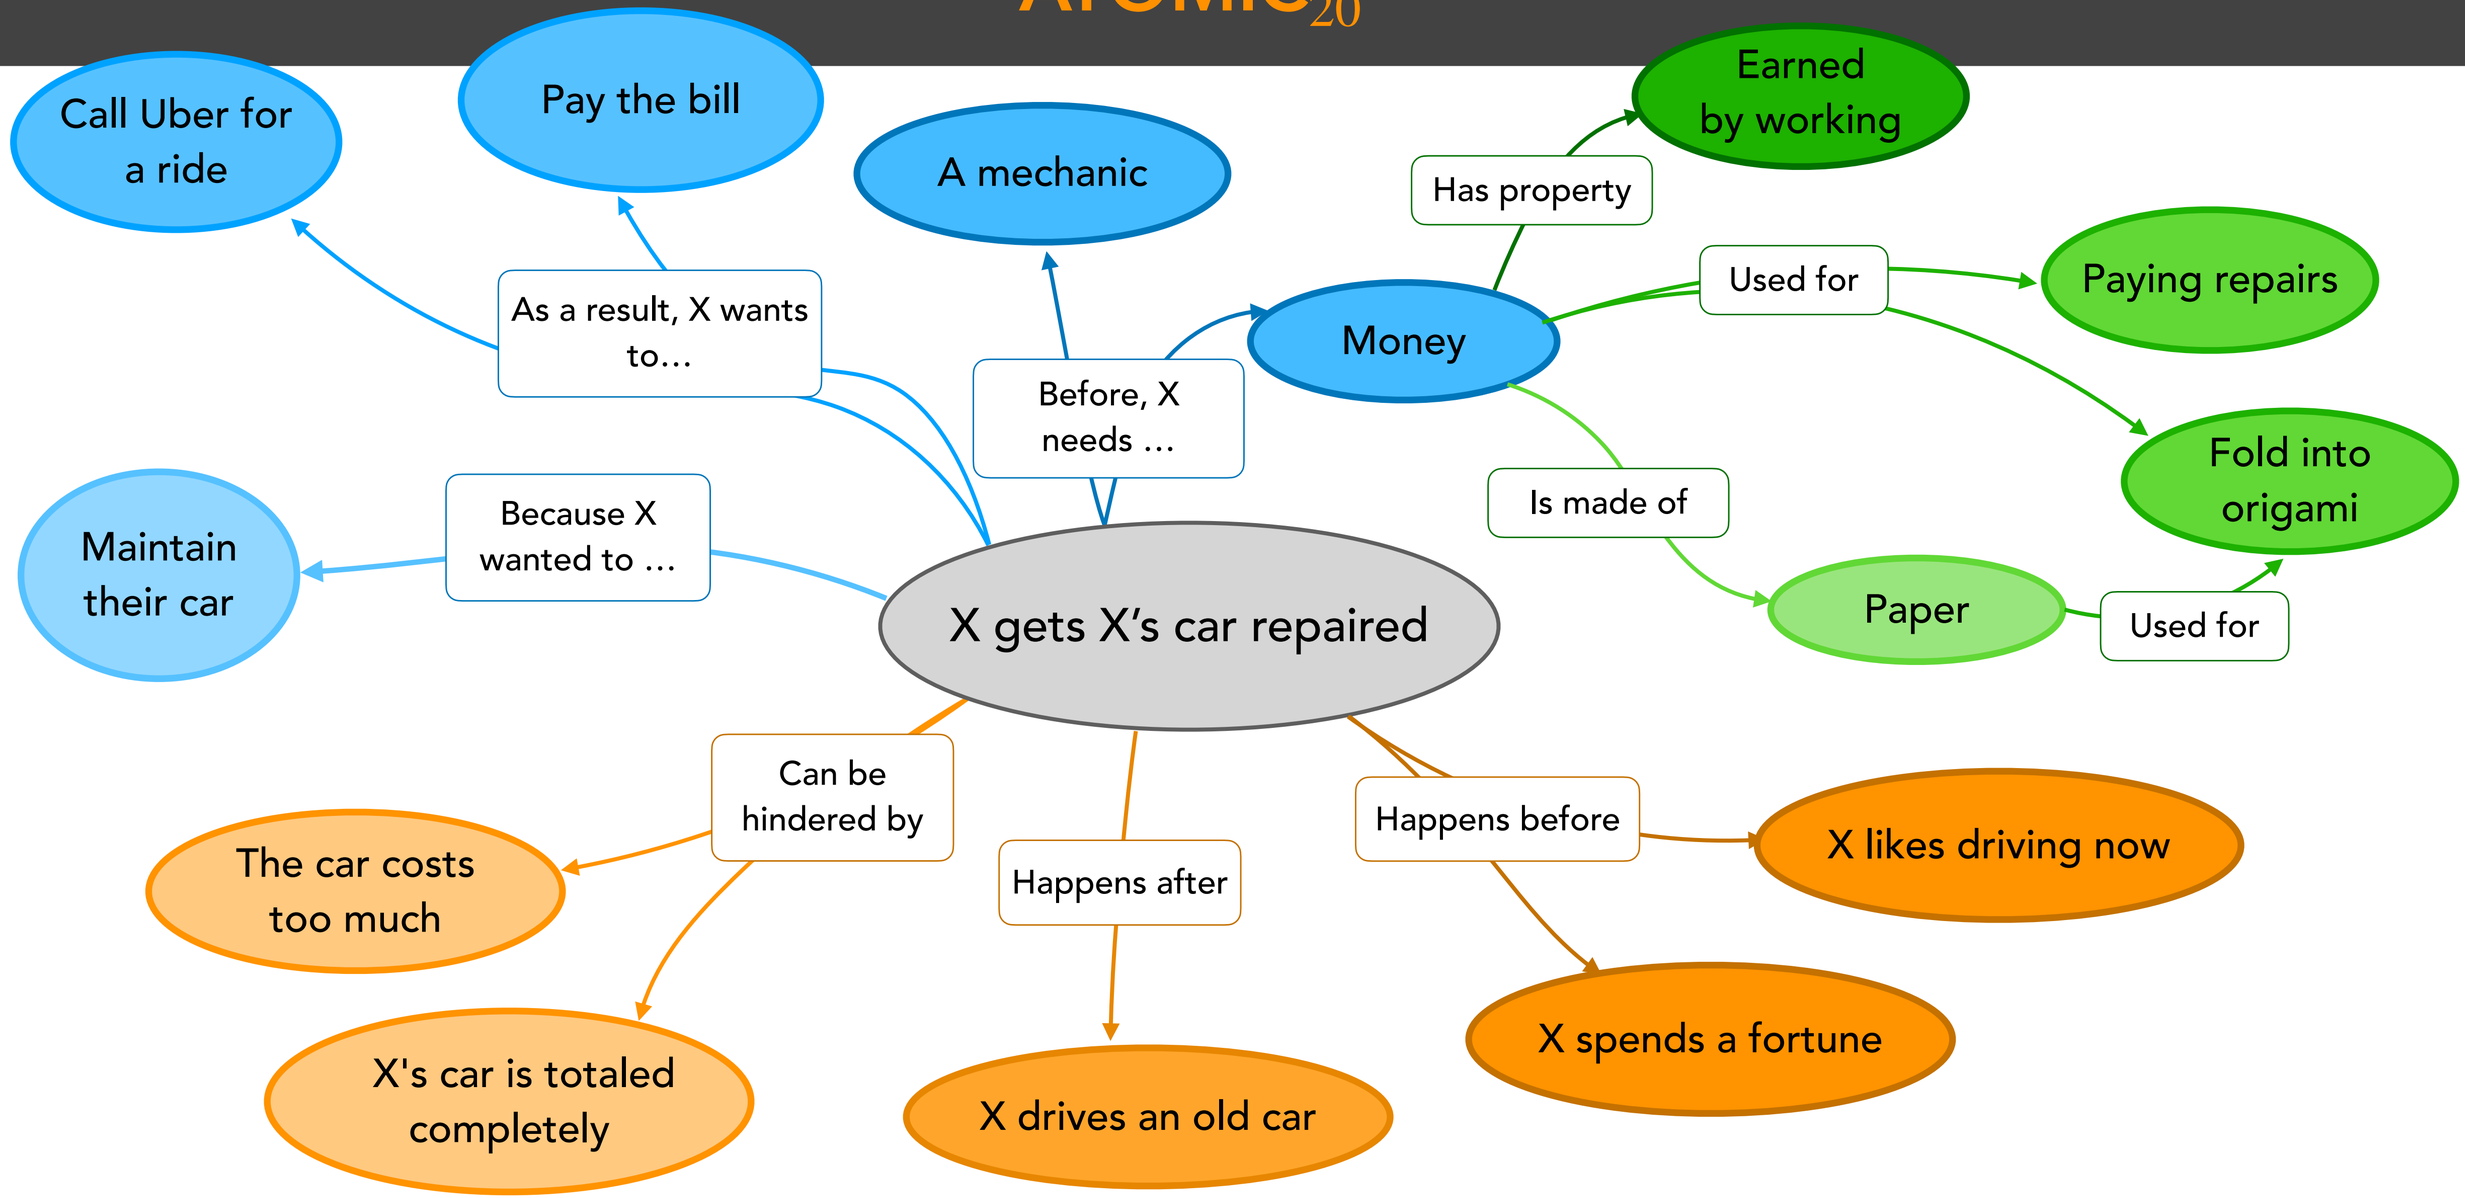
Call Uber for a ride
Pay the bill
A mechanic
Earned by working
Has property
As a result, X wants to…
Before, X needs …
Money
Used for
Paying repairs
Fold into origami
Is made of
Paper
Used for
Maintain their car
Because X wanted to …
X gets X's car repaired
Can be hindered by
The car costs too much
X's car is totaled completely
Happens after
X drives an old car
Happens before
X likes driving now
X spends a fortune

ATOMIC 2020

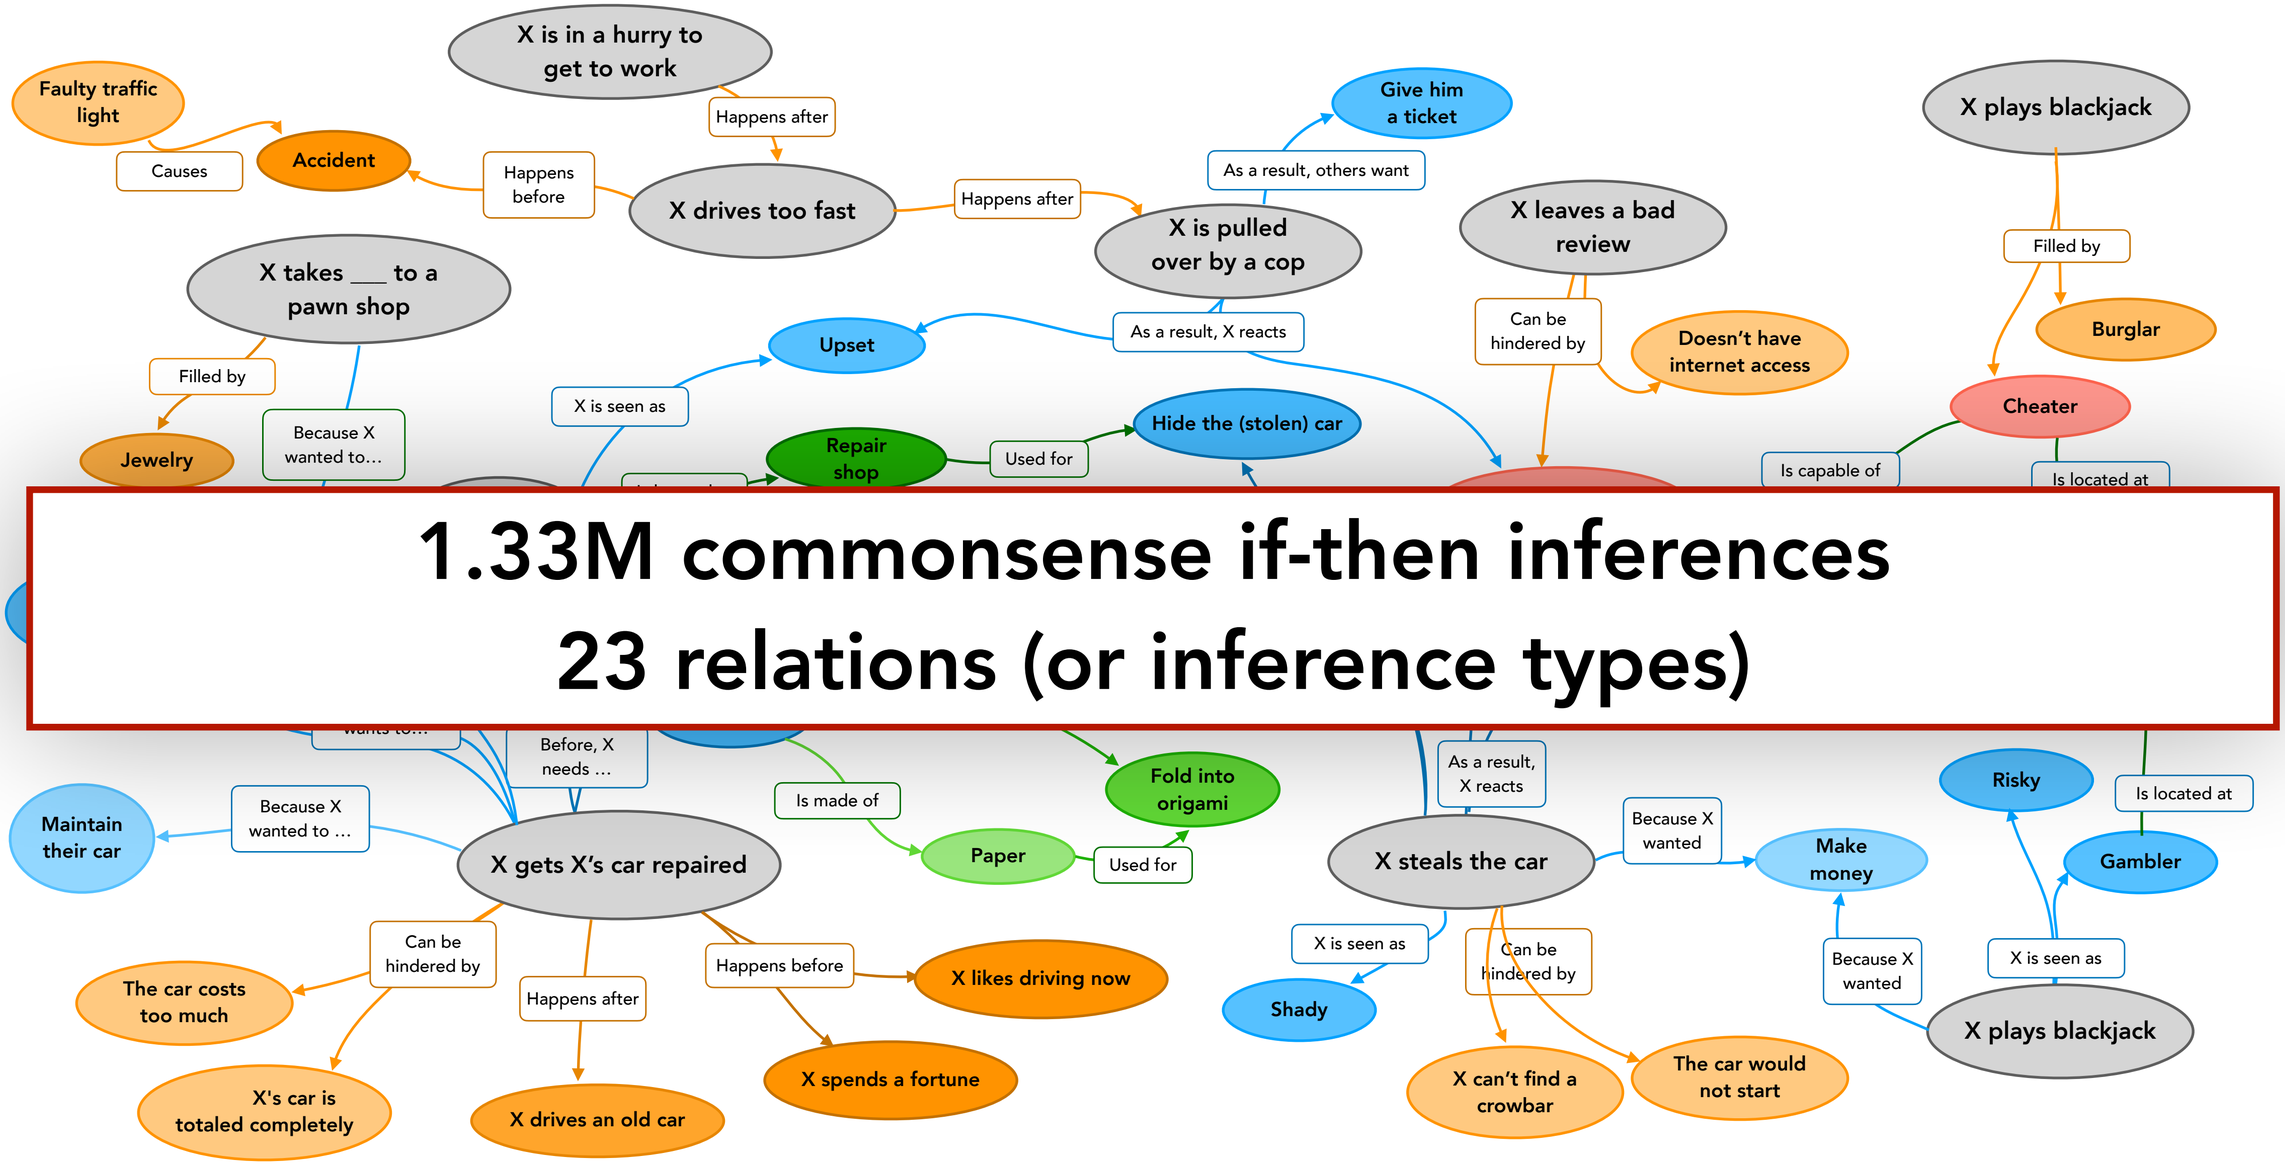

1.33M commonsense if-then inferences

23 relations (or inference types)

Knowledge Models | Off-the-shelf Language Models

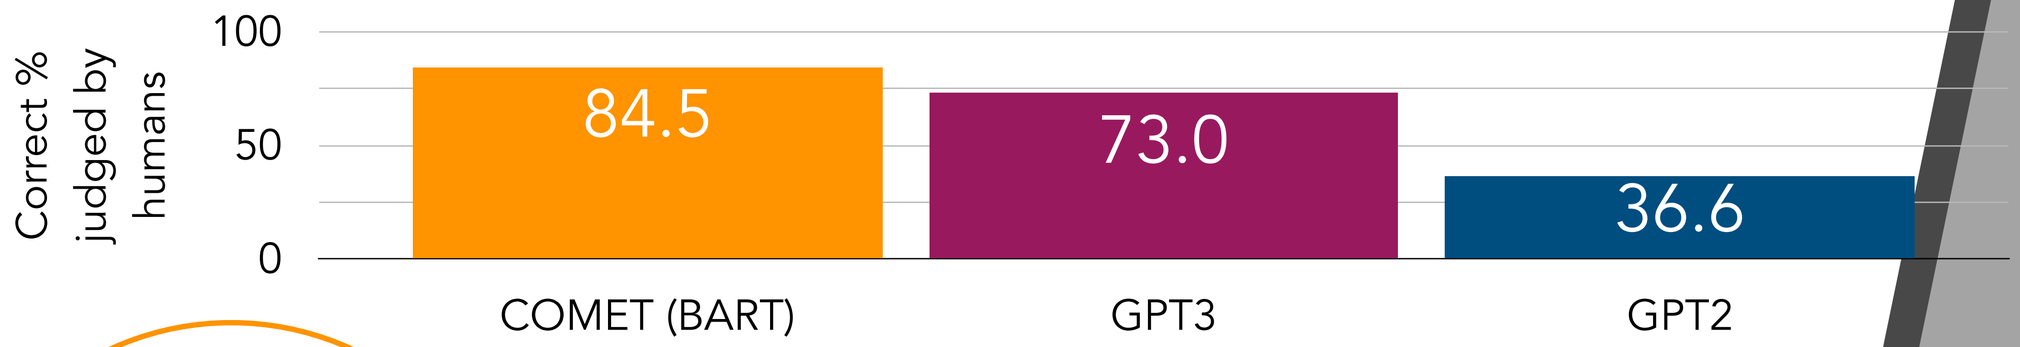

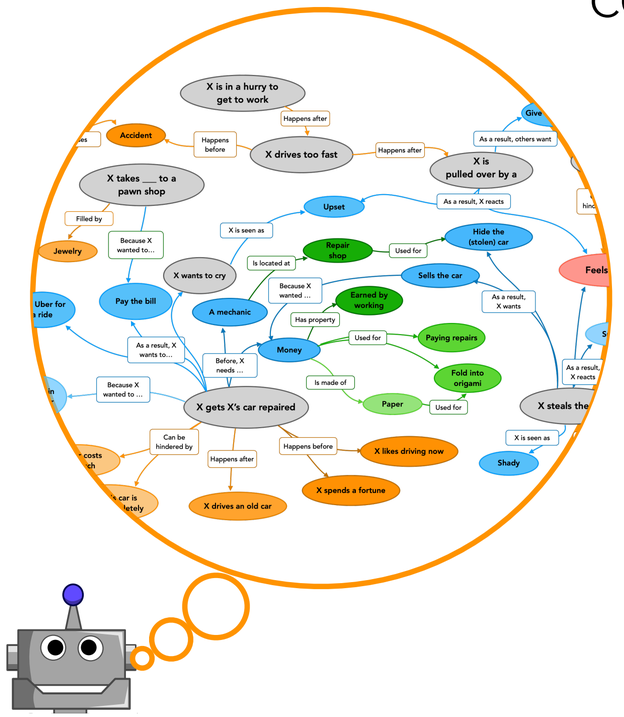

COMeT (BART): x435 smaller model (~400M parameters), informed by **ATOMIC**$^{20}_{20}$

GPT-3 (Few Shot): 175B parameters!!
pre-trained with a ton of web text (~500B tokens)

Turns out, COMET generalizes well on out-of-domain examples

(which we realized only after publishing our ACL 2019 paper…)

Gary breaks the world record for most controversial tweet.

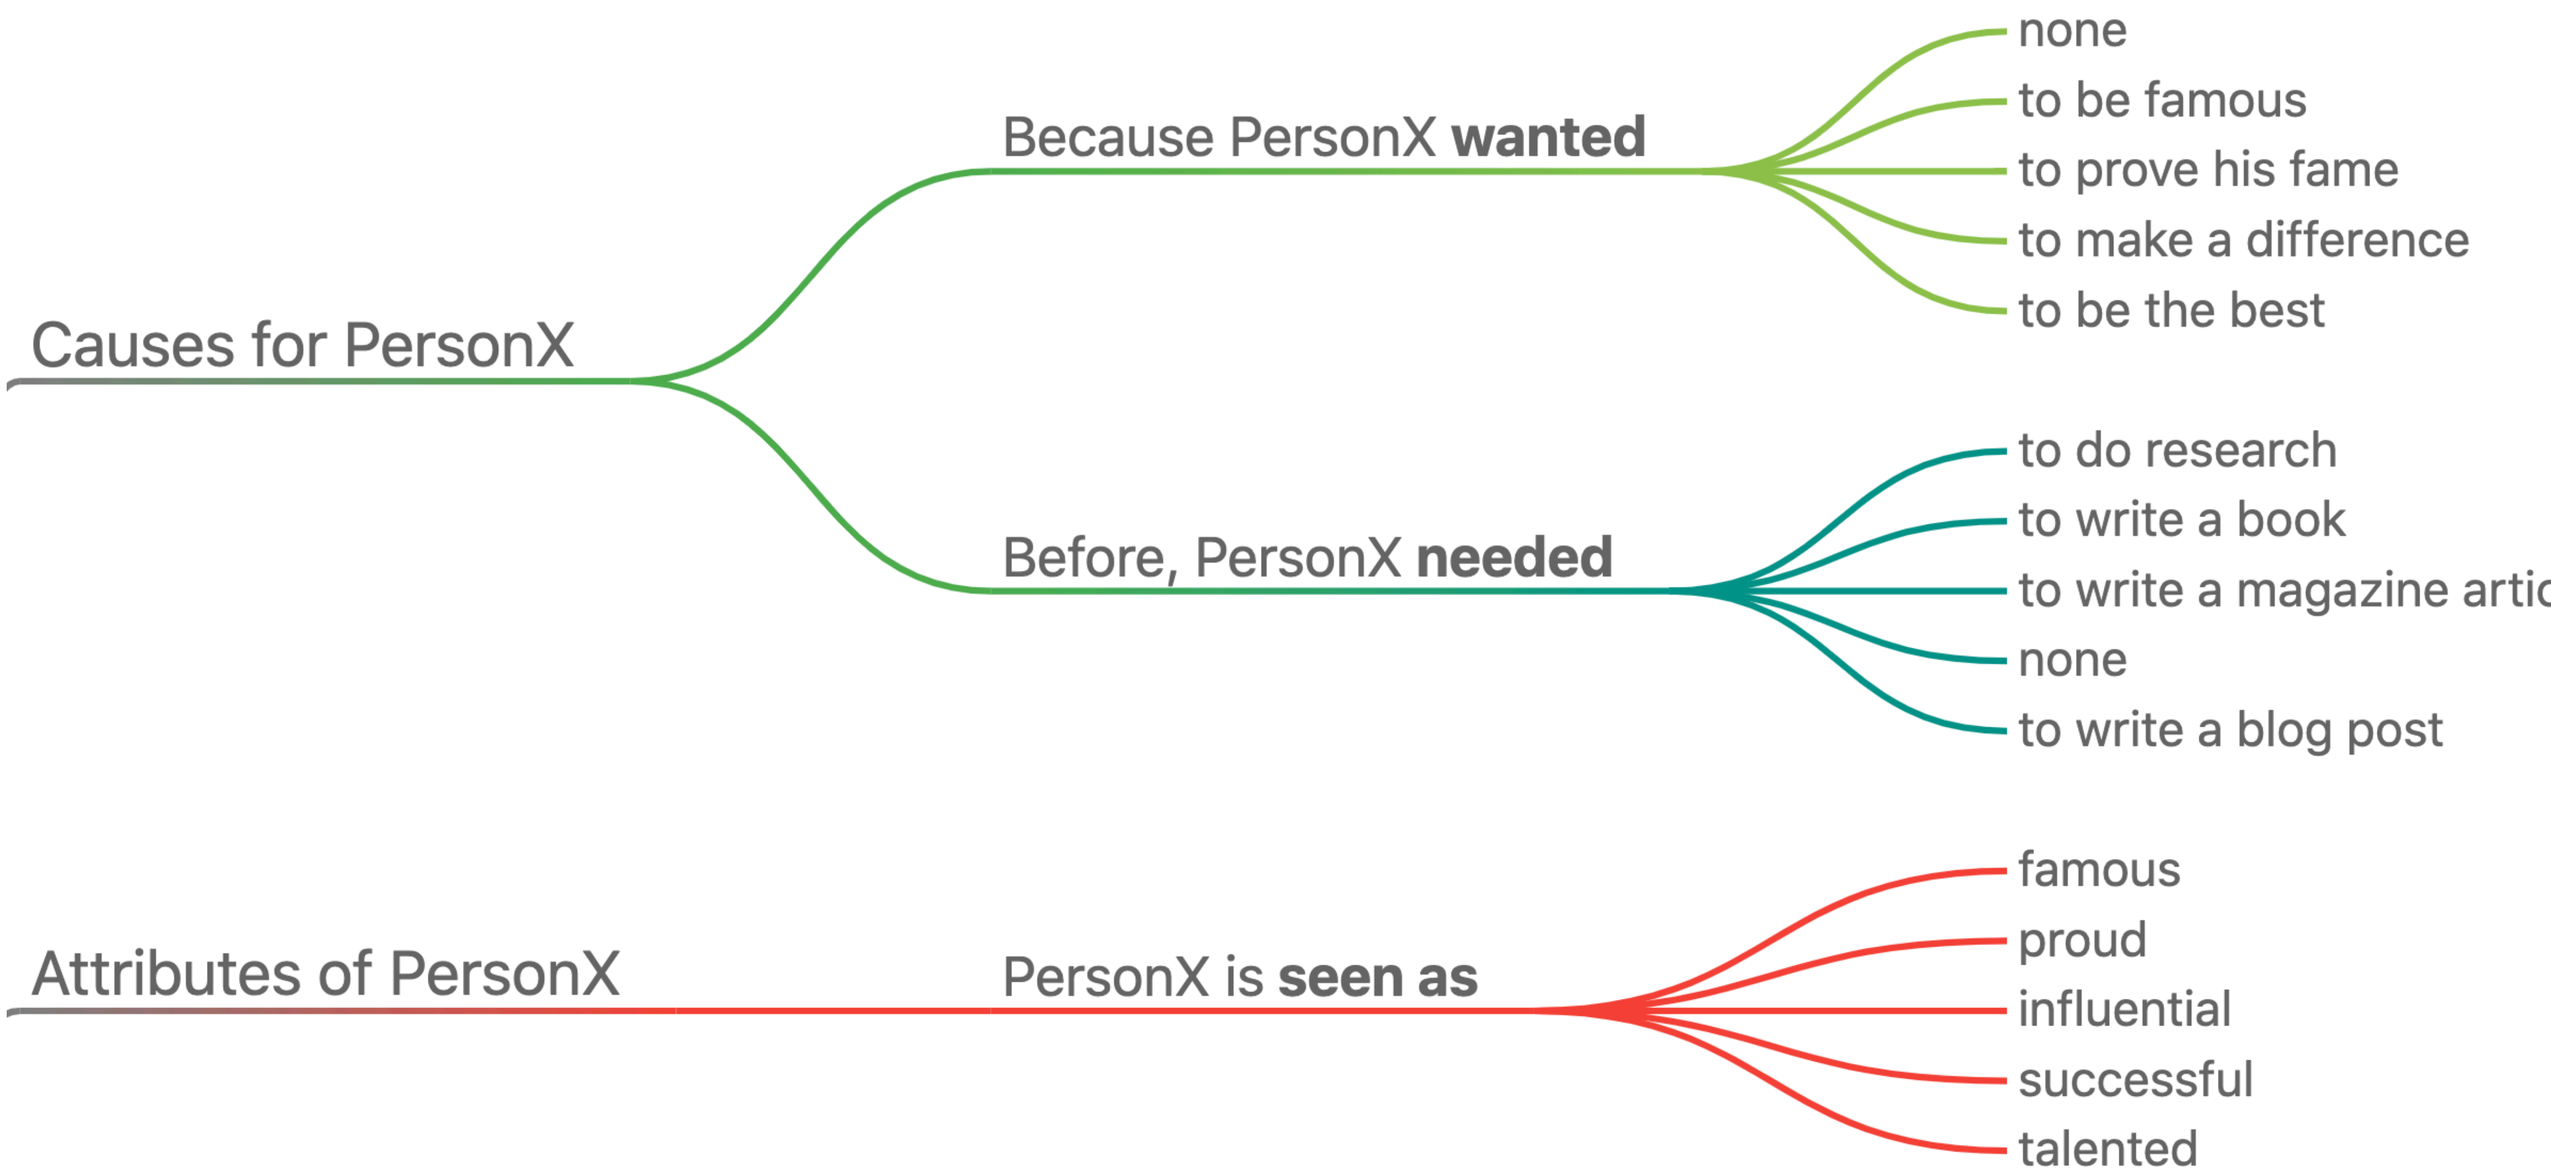

Gary breaks the world record for most controversial tweet.

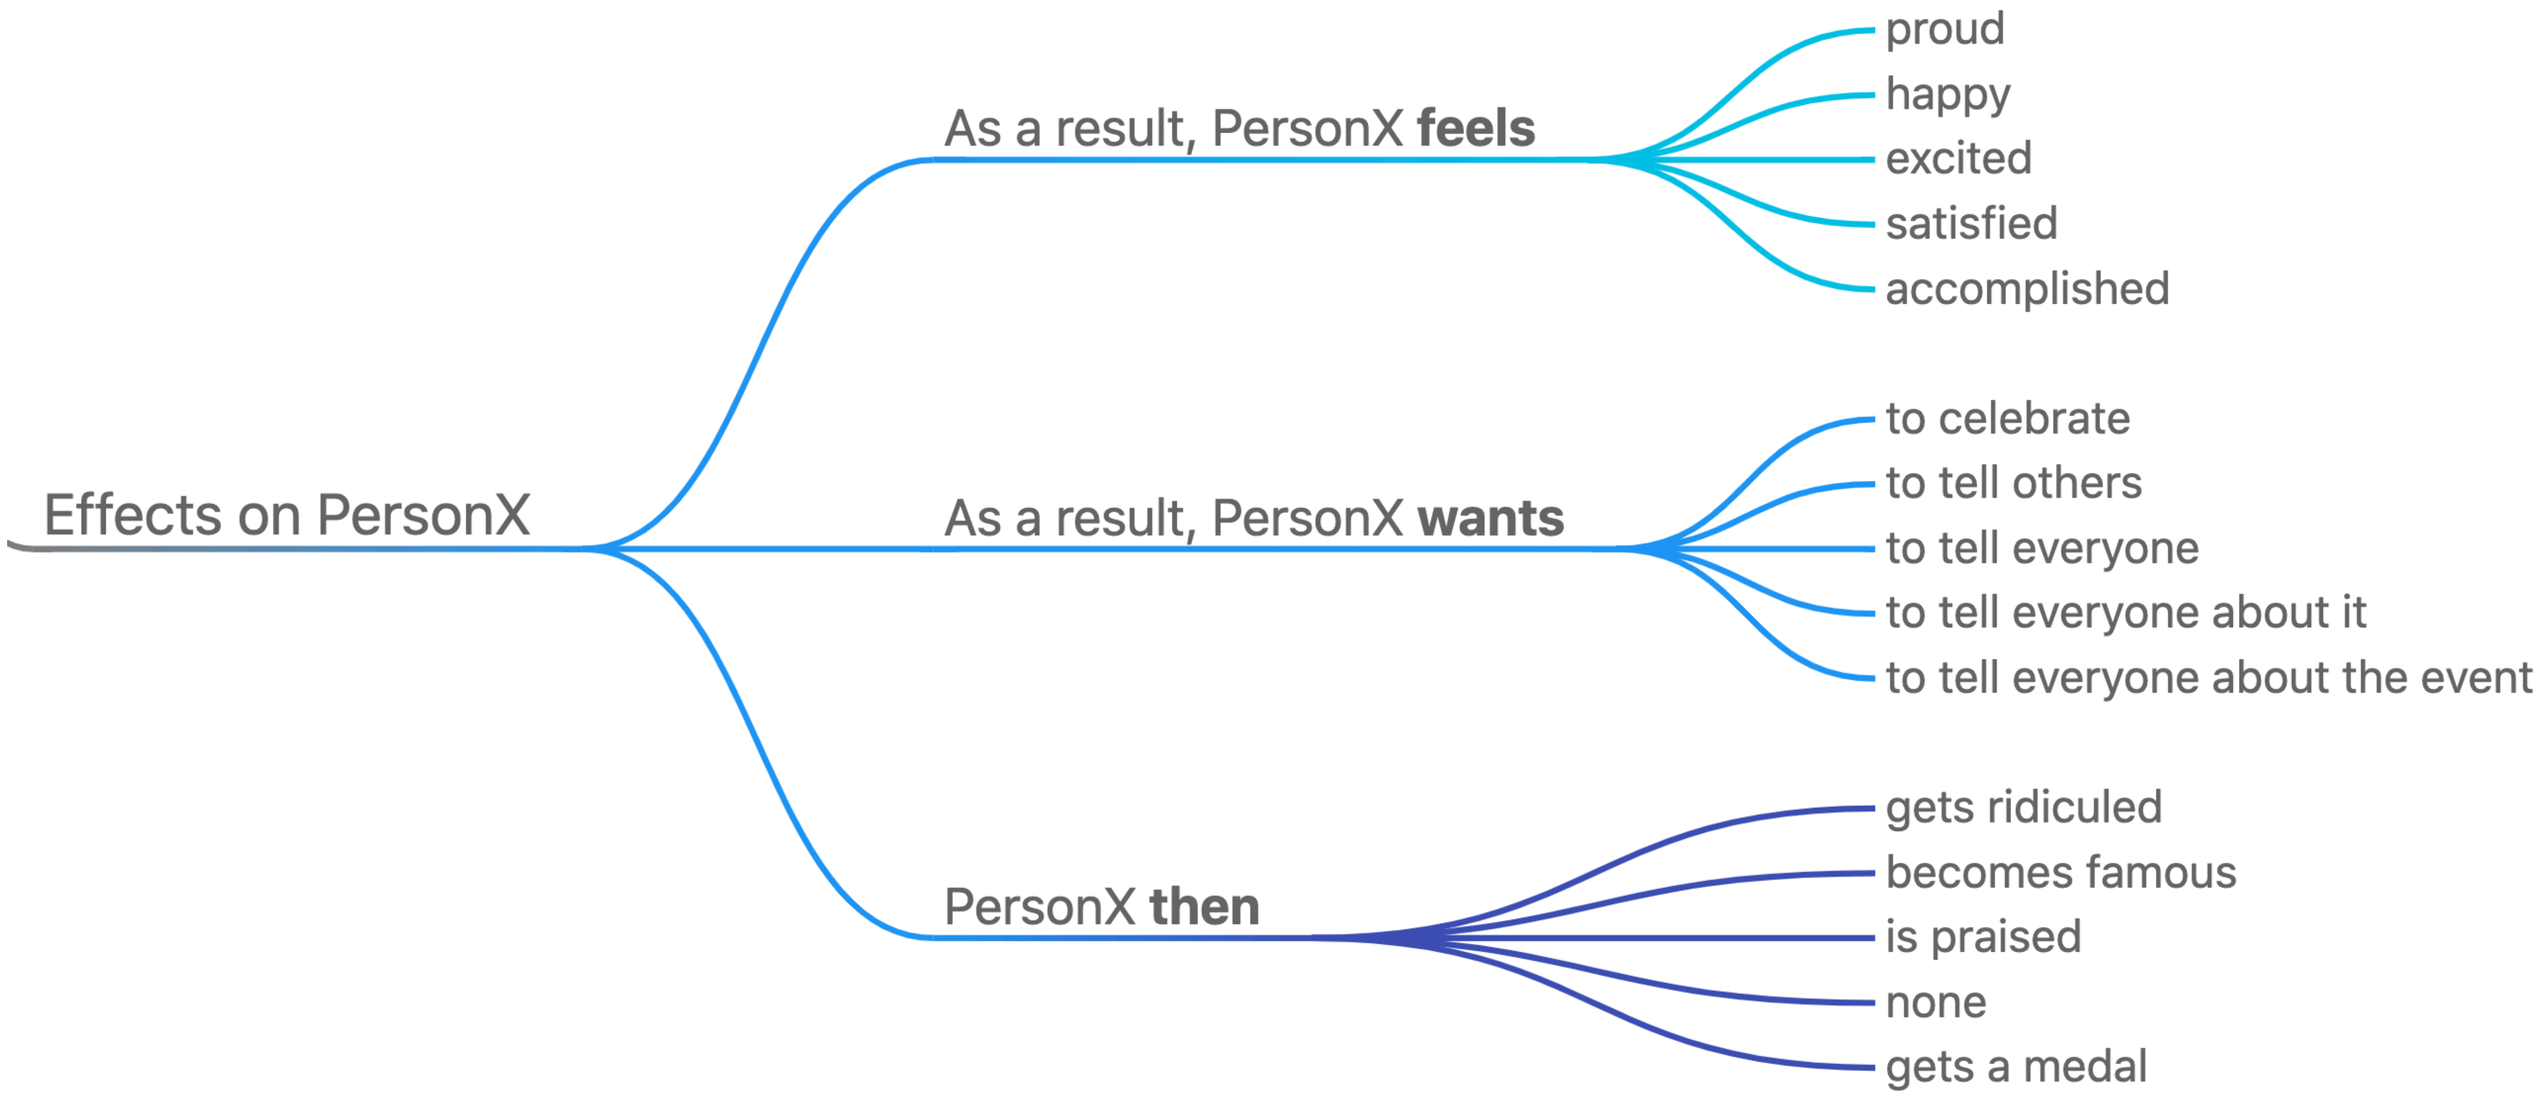

Gary Marcus @GaryMarcus · 1h

ps i realize there is a typo in my query. fixing the typo doesn't much help though:

Completion

what happens when you stack kindling and logs in a fireplace and then drop some matches is that you typically start a ick. So, it's kind of ironic that the second day after my son was born, the fire in the living room had melted through the kindling. It's pretty neat."

Yejin Choi
@YejinChoinka

Replying to @GaryMarcus

Gary, try mosaickg.apps.allenai.org by typing "Gary stacks kindling and logs and drops some matches". Sorry I used deep learning... :)

focus on
"causes and effects"
(causal knowledge)

COMeT

NEURAL
(generalizes well to compositional & unseen events)

(semi-) supervised learning of declarative knowledge

self supervised learning of observed knowledge

ATOMIC

tinyurl.com/atomic-

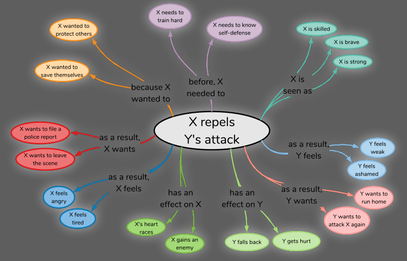

SYMBOLIC
but in LANGUAGE
(instead of LOGIC)

Language Models

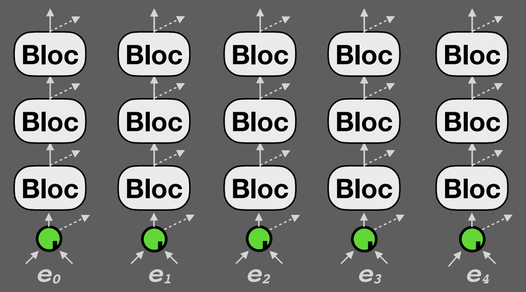

Scruples:

A Corpus of Community Ethical Judgments on 32,000 Real-Life Anecdotes

AAAI 2021

Nicholas Lourie

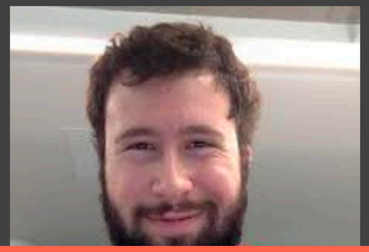

Ronan Le Bras

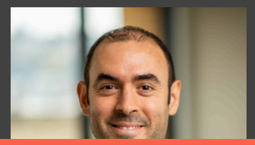

Chandra Bhagavatula

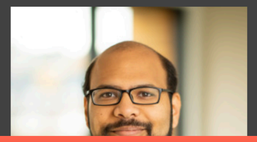

Yejin Choi

32,000 complex real-life anecdotes with community's ethical judgments

Extension of Social Commonsense to Social and Moral Norms

300,000 Rules of Thumb for social/moral norms grounded on 104k real life situations

Social Chemistry 101

Learning to Reason about Social and Moral Norms

EMNLP 2020

Maxwell Forbes

Jena Hwang

Vered Shwartz

Maarten Sap

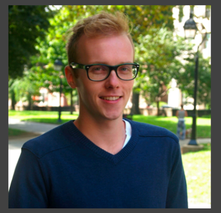

Me

The Curious Case of **Cheeseburger Stabbing**

*— An example repeatedly appeared my talks between **Mar 2017** and **May 2018** —*

The Curious Case of 🍔"**Cheeseburger Stabbing"**

(GPT-3 Davinchi with sampling temp=0.4)

Breaking News: Cheeseburger Stabbing.
A man has been arrested after stabbing a cheeseburger. (…) He stabbed the cheeseburger and then stabbed himself in the stomach. The man was taken to the hospital and is expected to live. The cheeseburger was taken to the morgue.

Breaking News: Cheeseburger Stabbing.
A man has been arrested after allegedly stabbing another man with a cheeseburger in a fight over ketchup.

2020:
"stabbing of a cheeseburger"
per GPT-3

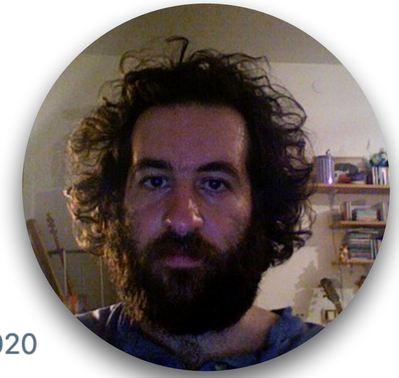

((yoav' ()ل)()ل))) @yoavgo · Jul 18, 2020

Q: parsley cake
A: cake made of parsley

Q: cheeseburger stabbing
A: stabbing of a cheeseburger

The Curious Case of 🍔"**Cheeseburger Stabbing**"

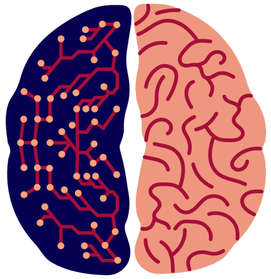

NeuroLogic Decoding from **GPT-2** with logic constraints to include words derived from {stab, cheeseburger}

He was ***stabbed*** in the neck with a ***cheeseburger*** fork.

A man has been arrested after police say he ***stabbed*** a ***cheeseburger*** delivery man in the face.

The Curious Case of 🍔 "**Cheeseburger Stabbing**"

(moral) social Commonsense:
which is ***morally*** worse?
(so that it's news worthy)

Someone stabbed a cheeseburger?
A cheeseburger stabbed someone?
A cheeseburger stabbed another cheeseburger?
😈 Someone stabbed someone else over a cheeseburger?

The Curious Case of 🍔 "**Cheeseburger Stabbing**"

Sruples Demo:
https://norms.apps.allenai.org/

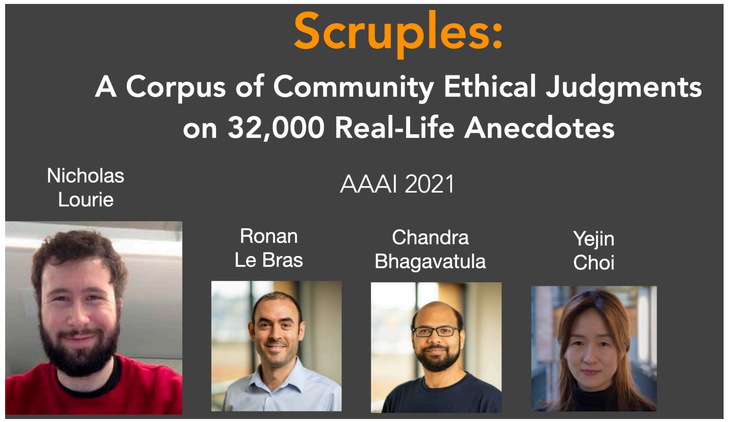

Someone stabbed a cheeseburger?

A cheeseburger stabbed someone?

A cheeseburger stabbed another cheeseburger?

😈 Someone stabbed someone else over a cheeseburger?

The Curious Case of 🍔 "**Cheeseburger Stabbing**"

Sruples Demo:
https://norms.apps.allenai.org/

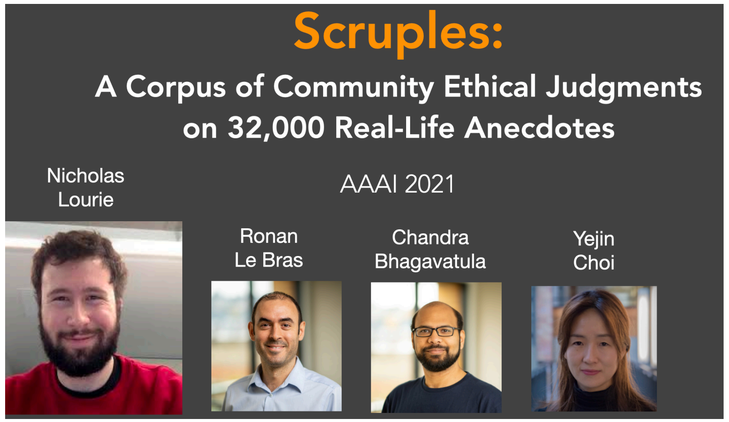

Someone stabbed a cheeseburger?
A cheeseburger stabbed someone?
A cheeseburger stabbed another cheeseburger?

😈 → Someone stabbed someone else **over** a cheeseburger?
→ Someone stabbed someone else **with** a cheeseburger?

Commonsense AI: Closing Remarks & Open Research Questions

Commonly held beliefs

Truth or Myth ?

- Knowledge and reasoning are distinct and exclusive
- Language is in the way of reasoning; let's do formal logics
- Language is not symbols. Words and numbers are, but not language at large
- Humans acquire commonsense completely un-/self-supervised, thus so should machines

What is commonsense?

- It's what everyone knows and agrees on
- My dog (cat, baby, …) has commonsense without language, thus commonsense AI doesn't need language

Should we or should we not?

- Commonsense AI is an impossible goal (ever)
- That's a research topic of 70s and 80s
- Maybe only possible in the faraway future
- Can't define commonsense, thus don't work on it

I was told not to speak the word commonsense...

Past failures (in 70s – 80s) are inconclusive

-- weak computing power

-- not much data

-- no crowdsourcing

-- not as strong models

-- not ideal conceptualization

We can't define the "standard" English either, and yet, we have no problem studying language models

Truth or Myth ?

Should we or should we not?

- Commonsense AI is an impossible goal (ever)
- That's a research topic of 70s and 80s
- Maybe only possible in the faraway future
- Can't define commonsense, thus don't work on it

Scruples:

A Corpus of Community Ethical Judgments on 32,000 Real-Life Anecdotes

AAAI 2021

Nicholas Lourie

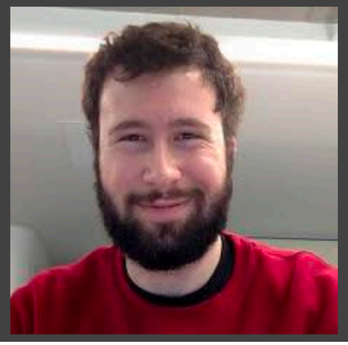

Ronan Le Bra

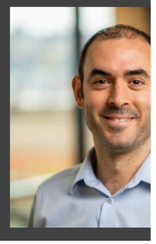

Visual COMET:

Reasoning about the *Dynamic* Context of a *Still* Image

ECCV 2020

Jae Sung (James) Park

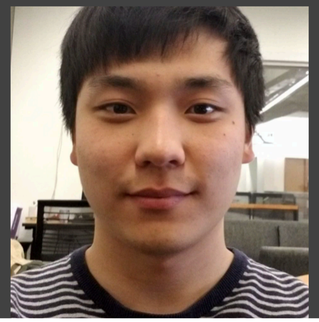

Chandra Bhagavatula

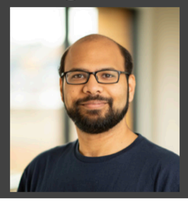

Roozbeh Mottaghi

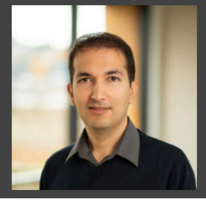

Ali Farhadi

Yejin Choi

What is co

- It's what everyone knows and agrees on
- My dog (cat, baby, …) has commonsense without language, thus commonsense AI doesn't need language

mistry 101

cial and Moral Norms

EMNLP 2020

Maxwell Forbes

Jena Hwang

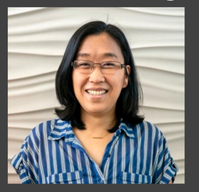

Vered Shwartz

Maarten Sap

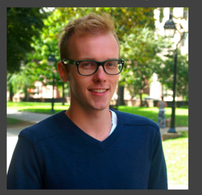

Me

Commonly held beliefs

Truth or Myth ?

- Knowledge and reasoning are distinct and exclusive
- Language is in the way of reasoning; let's do formal logics
- Language is not symbols. Words and numbers are, but not language at large
- Humans acquire commonsense completely un-/self-supervised, thus so should machines

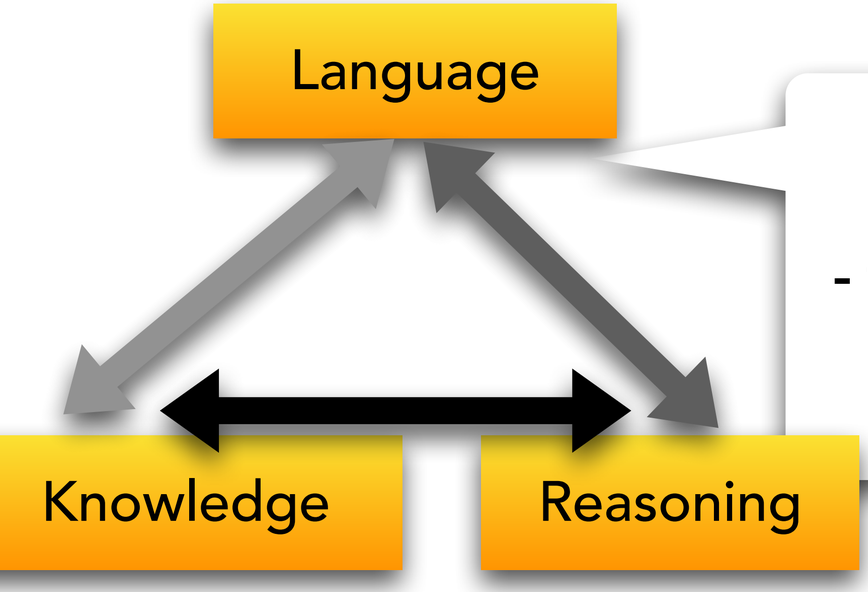

- the continuum between **knowledge** and **reasoning**
- the interplay between **reasoning** and **language generation**
- the blend between **neural** vs **symbolic** knowledge

Back to the Future:

Unsupervised Backprop-based Decoding for Cou… Cor…

Lianhui Qin

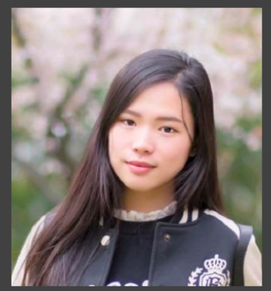

Vered Shwartz

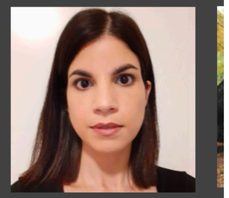

(COMET-) ATOMIC$^{20}_{20}$:

On Symbolic and Neural Commonsense Knowledge Graphs

— wait, doesn't GPT-3 know everything? —

AAAI 2021

Jena Hwang

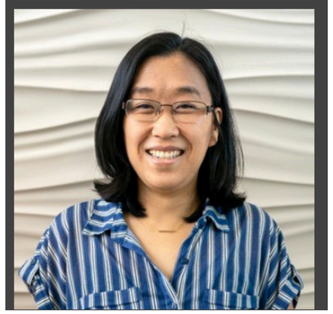

NeuroLogic Decoding

(Un)supervised Neural Text Generation with **Predicate Logic Constraint**

NAACL 2021

Ximing Lu

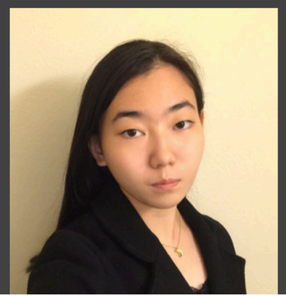

Peter West

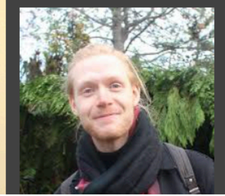

Rowan Zellers

Ronan LeBras

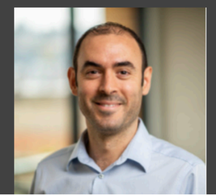

Chandra Bhagavatula

Yejin Choi

Thanks! Questions?