Commonsense Reasoning

Lara J. Martin (she/they)

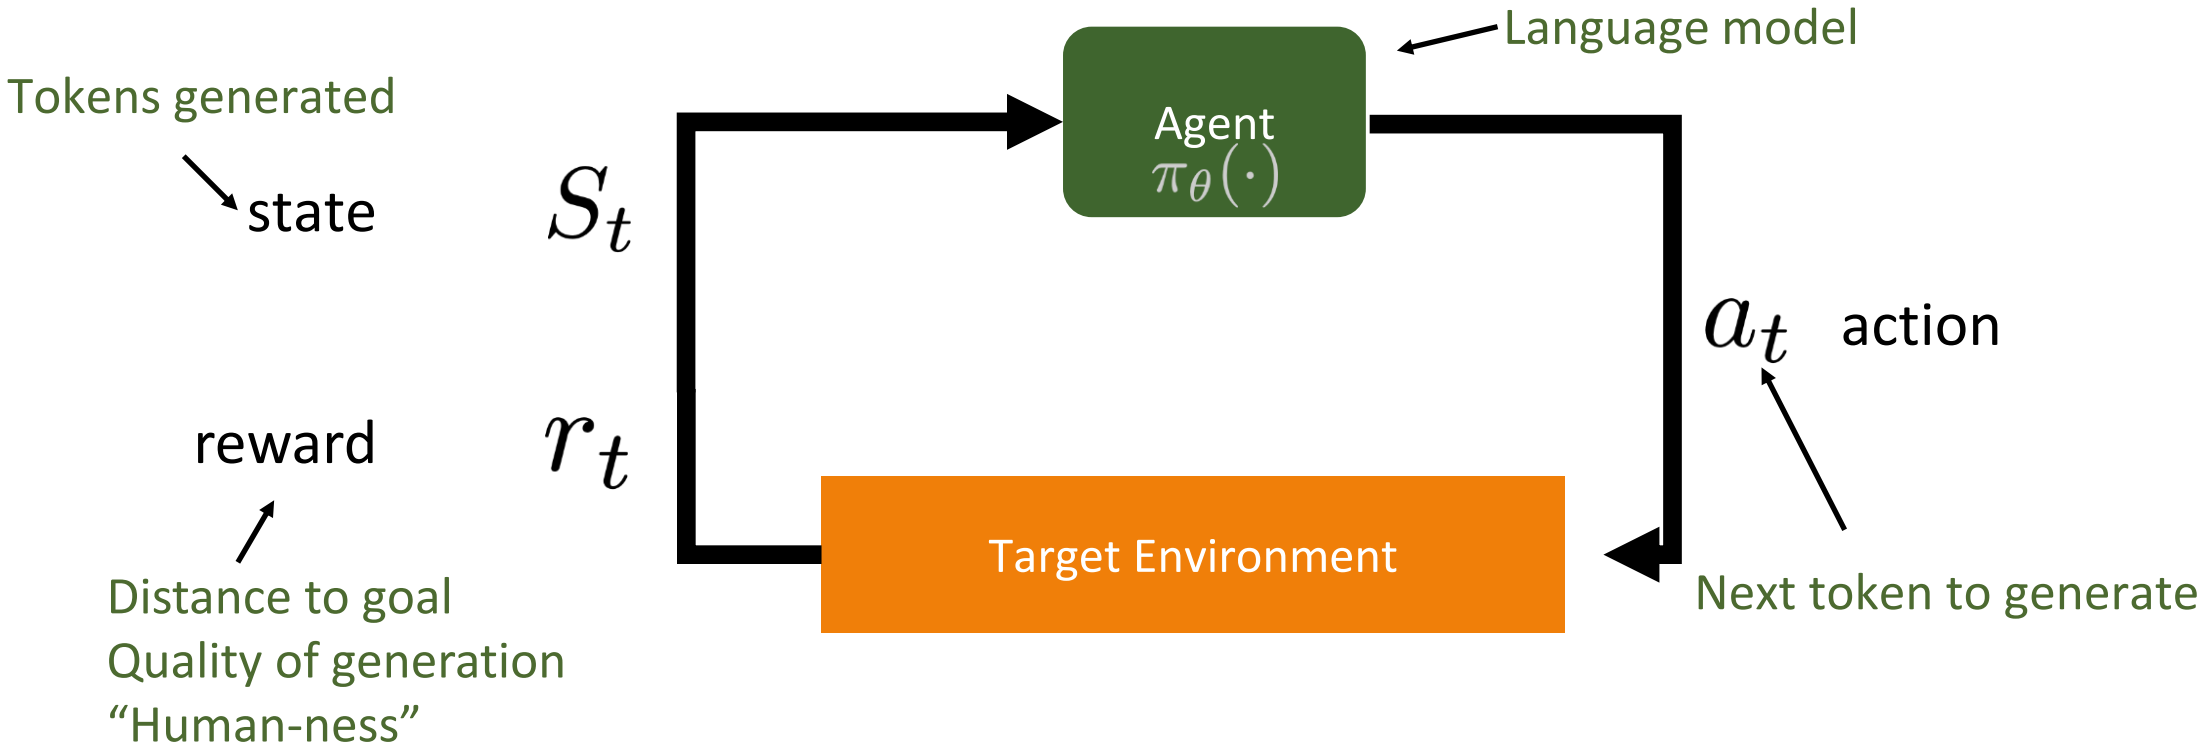
https://laramartin.net/neurosymbolic-text-gen/

Modified from slides by Chris Callison-Burch

Learning Objectives

- Recognize the similarities and differences between human commonsense and machine reasoning
- Consider the state of machine intelligence and reasoning
- Categorize types of commonsense
- Find out about existing popular commonsense knowledge bases
- Distinguish when you would use one knowledge base over another

Review:
RL in the Context of Language Models…



$a_t \sim \pi_\theta(S_t)$: policy

Review: REINFORCE

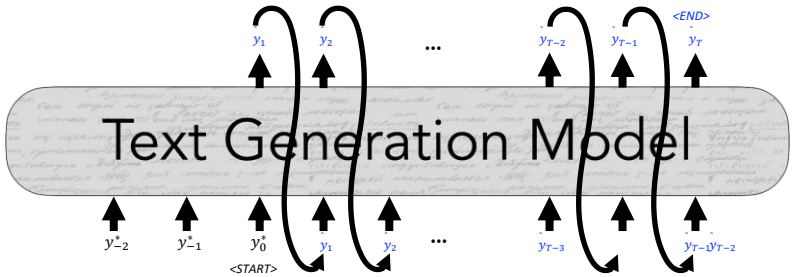
- Sample a sequence from your model, score the sequence, and use the score to train the model.

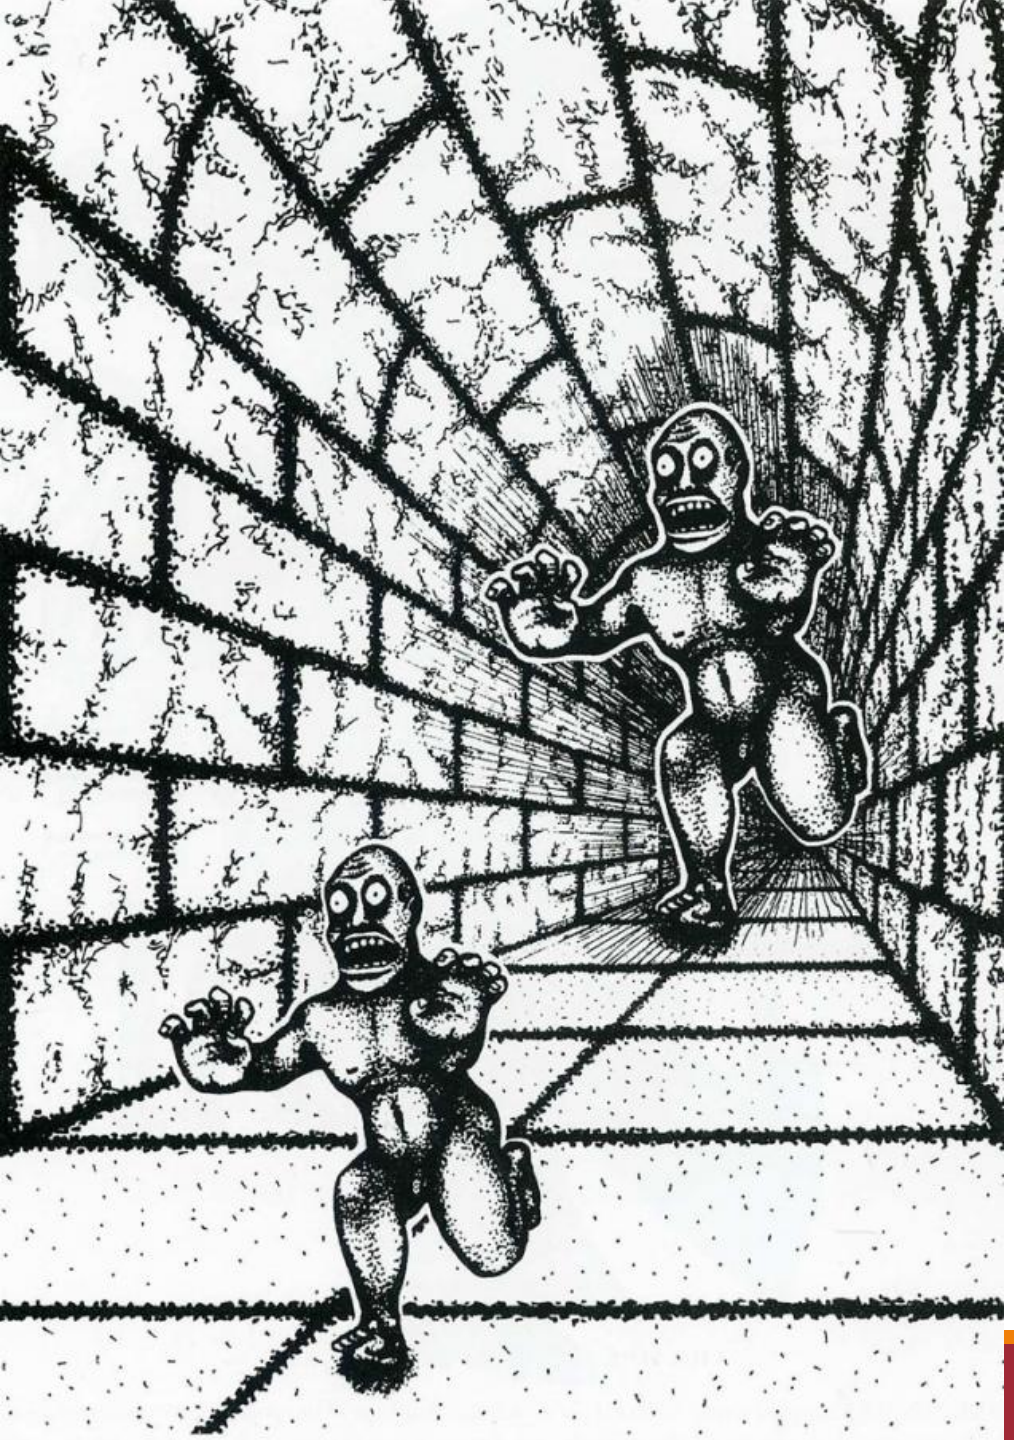
Next time, increase the probability of this sampled token in the same context.

$$L_{RL} = -\sum_{t=1}^{T} r(\hat{y}_t) \log P(\hat{y}_t | \{y^*\}; \{\hat{y}\}_{<t})$$

... but increase it more if I get a higher reward from the reward function.

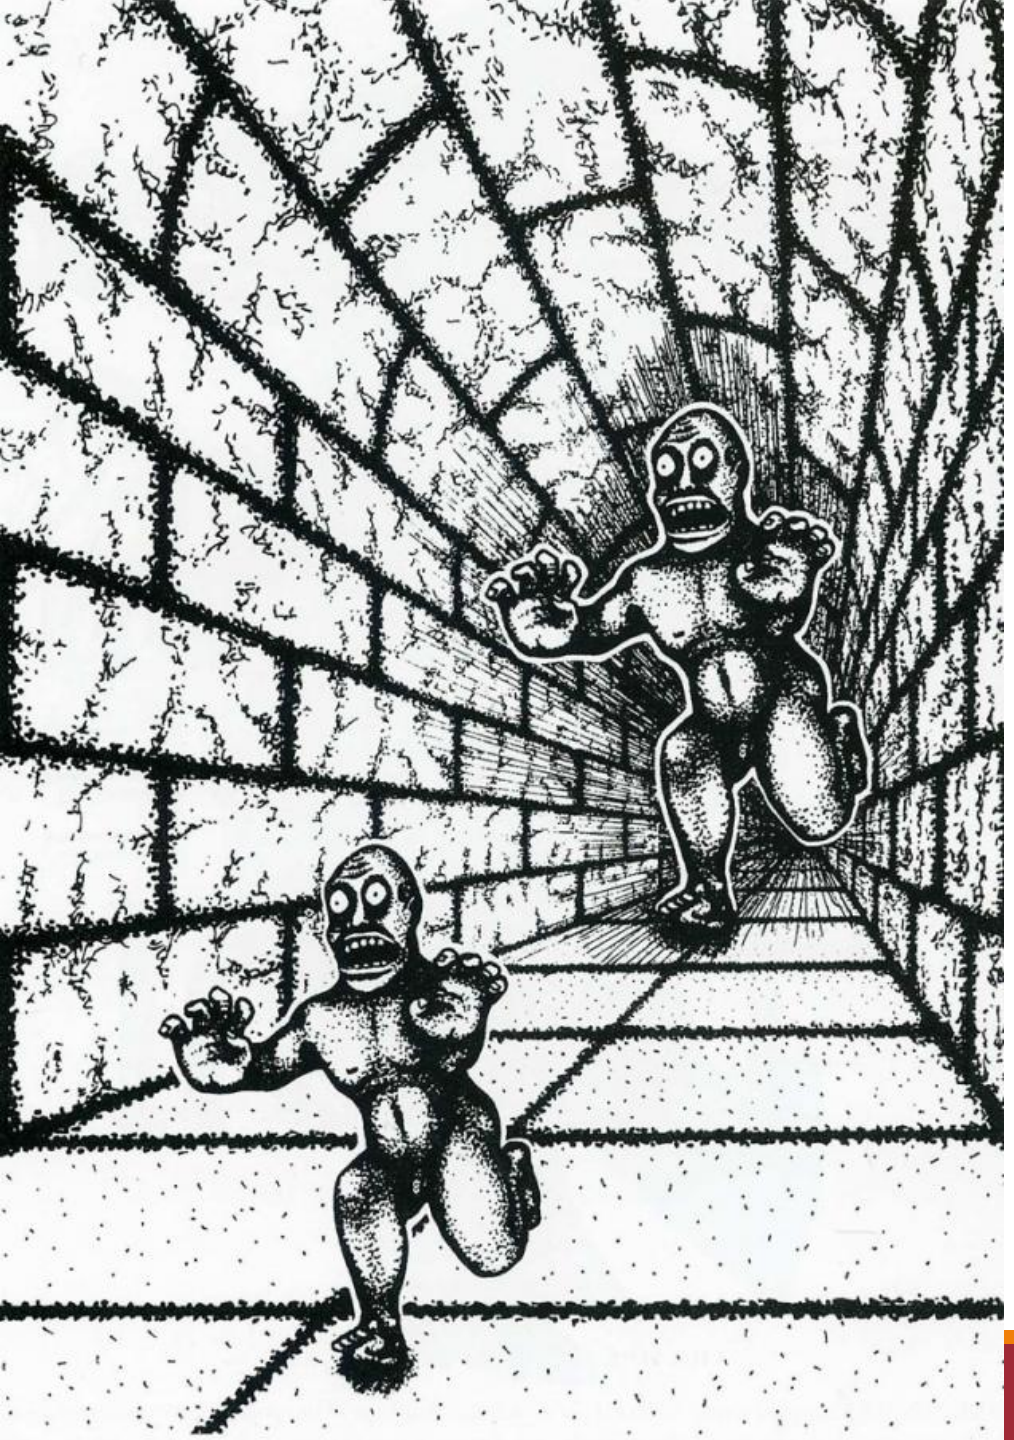
- $r(\cdot)$: Your reward model
- y^* :Input sequence given to the model
- $\hat{y}$:The sequence sampled from the model given y^*





What does this picture show?

Image from Roger Shepard, “Terror Subterra” from *Mind Sights: Original Illusions, Ambiguities, and other Anomalies* (1990)



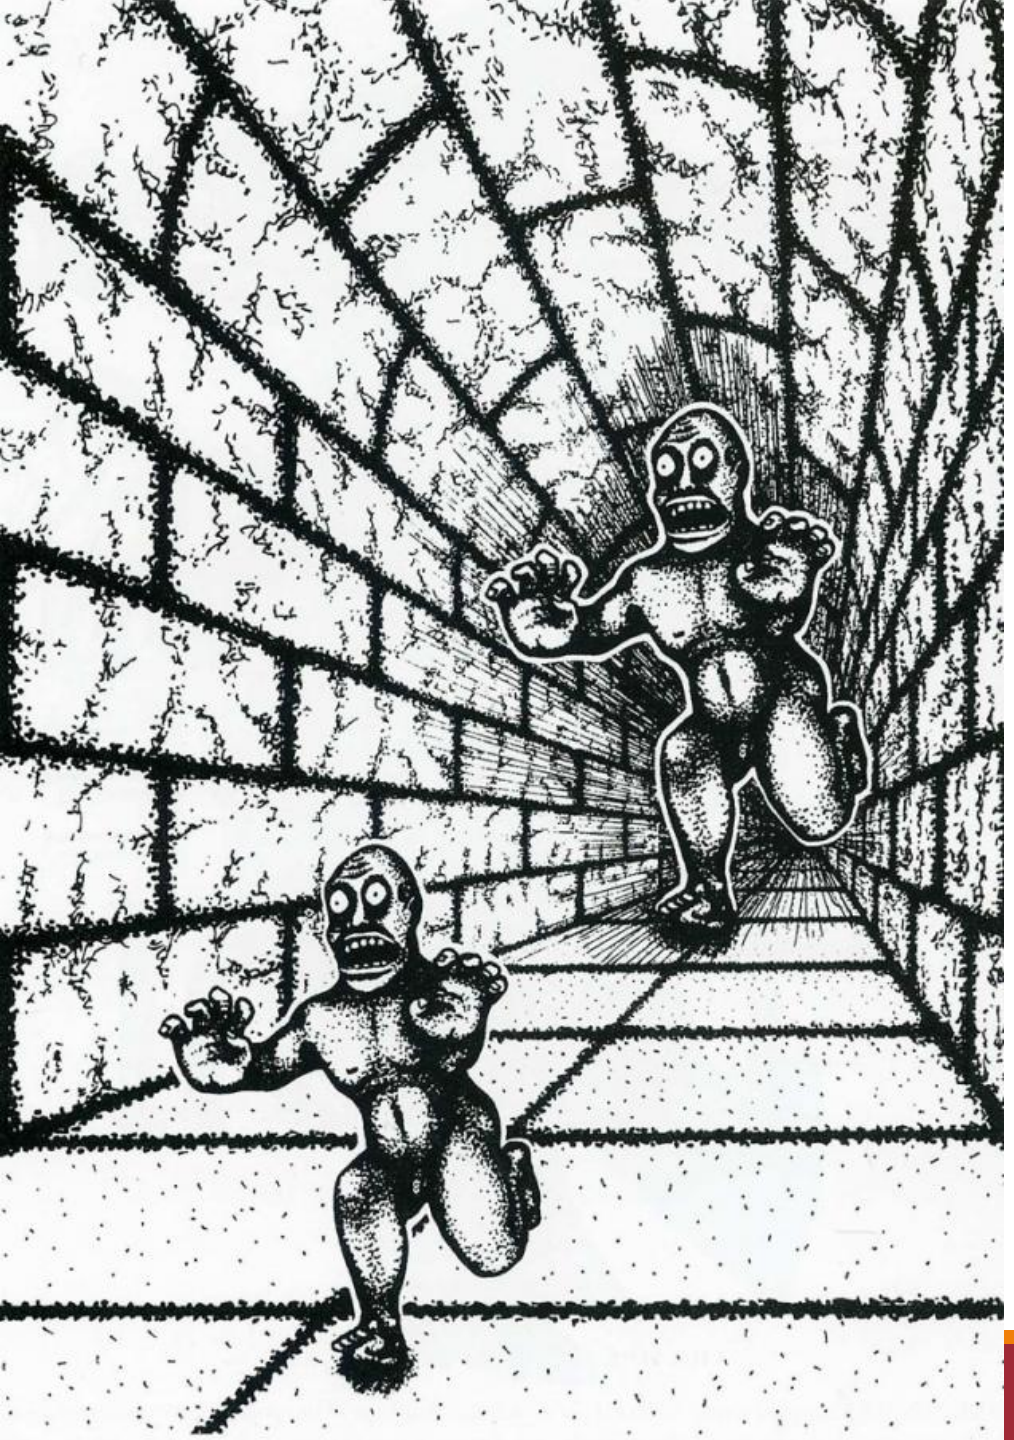
Monsters in a Tunnel

- **Two monsters are running** (rather than standing still on one foot)
- **One is chasing another** (rather than trying to copy his movements)
- **The chaser has hostile intentions and the chased is afraid** (even though two faces are identical)

Important Observations:

- A great deal of **intuitive inferences** are **commonsense inferences**, which can be described in **natural language**.
- None of these inferences is absolutely true. The inferences are **stochastic** in nature. Everything is **defeasible** with additional context.
- Commonsense inferences are about **predicting new information** that **is likely to be true** based on partially available information.

Image from Roger Shepard, "Terror Subterra" from *Mind Sights: Original Illusions, Ambiguities, and other Anomalies* (1990)



Powerpoint provides this alt text: “Black-and-white comic-style illustration depicts two startled, blob-like humanoid figures running through a long perspective corridor. Dense brick-like wall lines, converging floor and ceiling lines, and exaggerated facial expressions create a surreal sense of depth and motion.”

How do you feel about this description?

Image from Roger Shepard, “Terror Subterra” from *Mind Sights: Original Illusions, Ambiguities, and other Anomalies* (1990)

Claims of AI systems reaching a “human level”

INNOVATION

Experts Explain What AI Superintelligence Means And How Close We Are

By Joe McKendrick, Senior Contributor. Edited by Editor.
for Technology

Published Oct 20, 2026, 06:30am EDT

OpenAI says it cracked 90-year-old maths problem in 88 hours

8 September 2026

Share Save Add as preferred on Google

Kali Hays

Forbes

AI Claims Progress for 372 Hard Math Problems. The Competitive Impact

Paul Baier, Contributor
October 14, 2026 • 2 min read

Artificial Intelligence

AI Has Quietly Surpassed Human Intelligence, Are We Ready to Admit It?

Sam Altman says AI takeoff has begun, machines may already outthink humans

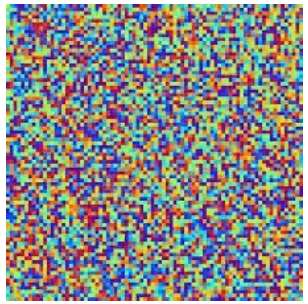
AI isn’t conscious—but we may be bringing it to life

In rethinking whether AI is sentient, we are asking bigger questions about cognition, human-machine interaction and even our own consciousness

BY SIMON DUAN



+



=



Giant panda

Gibbon

Object Recognition

Szegedy et al, 2014….



VQA

Jabri et al, 2017

Older methods may have ”solved” datasets rather than the underlying “task”



A horse standing in the grass.

Captioning

MacLeod et al, 2017

…a moved to Prague in 1880. … Tadakatsu moved to Chicago in 1881.

Where did Tesla move in 1880? Chicago

QA

Jia et al, 2017

AI solved one of math's hardest problems. Humanity learned nothing (so far)

UPDATED SEPTEMBER 22, 2026 · 7:08 AM ET

HEARD ON ALL THINGS CONSIDERED



Geoff Brumfiel

"So far it's been very difficult to really extract any human understanding from this new AI proof," said James Maynard, a mathematician at the University of Oxford.

https://www.npr.org/2026/09/22/nx-s1-5968588/openai-navier-stokes-problem-mathematicians-learn-little

OpenAI withdraws three preprints a day after releasing 722 manuscripts on unsolved math problems

The Association for Human Mathematics criticized the most recent release as well. “Mathematicians did not ask for this work to be done,” the group’s statement, released Oct. 7, said. “Releasing over 700 files at once is not a demonstration of scholarship, but a demonstration of power. We urge mathematicians to discontinue their work with OpenAI and to return to a vision of science that centers human understanding.”

https://retractionwatch.com/2026/10/08/openai-withdraws-preprints-722-manuscripts-unsolved-math-problems/

What is intelligence?

“the ability to learn or understand things or to deal with new or difficult situations” (Merriam-Webster)

Think-Pair-Share

Is reasoning solved? Why or why not?

Intelligence Testing

Modern intelligence testing was developed in 1905 to “select children for placement in special education classes” [1]

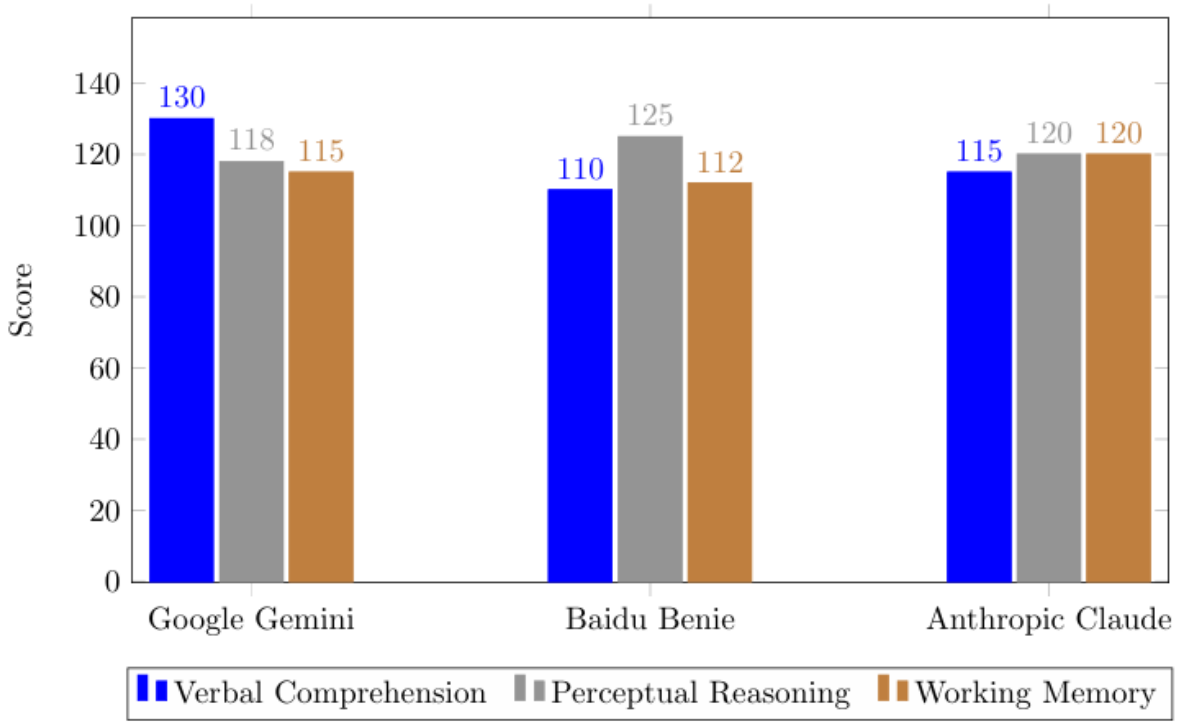


Fig. 1 Average Scores of LLMs on Adapted WAIS Sections

[1] WHITE, MARGARET B., and ALFRED E. HALL. “An Overview of Intelligence Testing.” *Educational Horizons*, vol. 58, no. 4, 1980, pp. 210–16. *JSTOR*, http://www.jstor.org/stable/42924403.

Image: Eryk Wasilewski, Mirek Jablonski. Measuring the Perceived IQ of Multimodal Large Language Models Using Standardized IQ Tests. *TechRxiv*. 13 May 2024. DOI: https://doi.org/10.36227/techrxiv.171560572.29045385/v1

What is commonsense?

Theory of Core Knowledge

Domain	Description
Objects	supports reasoning about objects and the laws of physics that govern them
Agents	supports reasoning about agents that act autonomously to pursue goals
Places	supports navigation and spatial reasoning around an environment
Number	supports reasoning about quality and how many things are present
Forms	supports representation of shapes and their affordances
Social Beings	supports reasoning about Theory of Mind and social interaction

Developmental psychologists have shown that children develop the ability to reason about these domains early in life. Such reasoning is important for later learning.

Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. Developmental Science, 10(1), 89-96. doi: 10.1111/j.1467-7687.2007.00569.x

Definition of Common Sense

The basic level of **practical knowledge** and **reasoning**

concerning **everyday situations and events**

that are **commonly shared** among most people.

It's OK to keep the closet door open

It's not OK to keep the refrigerator door open because the food might go bad

Essential for humans to live and interact with each other in a reasonable and safe way

Essential for AI to understand human needs and actions better

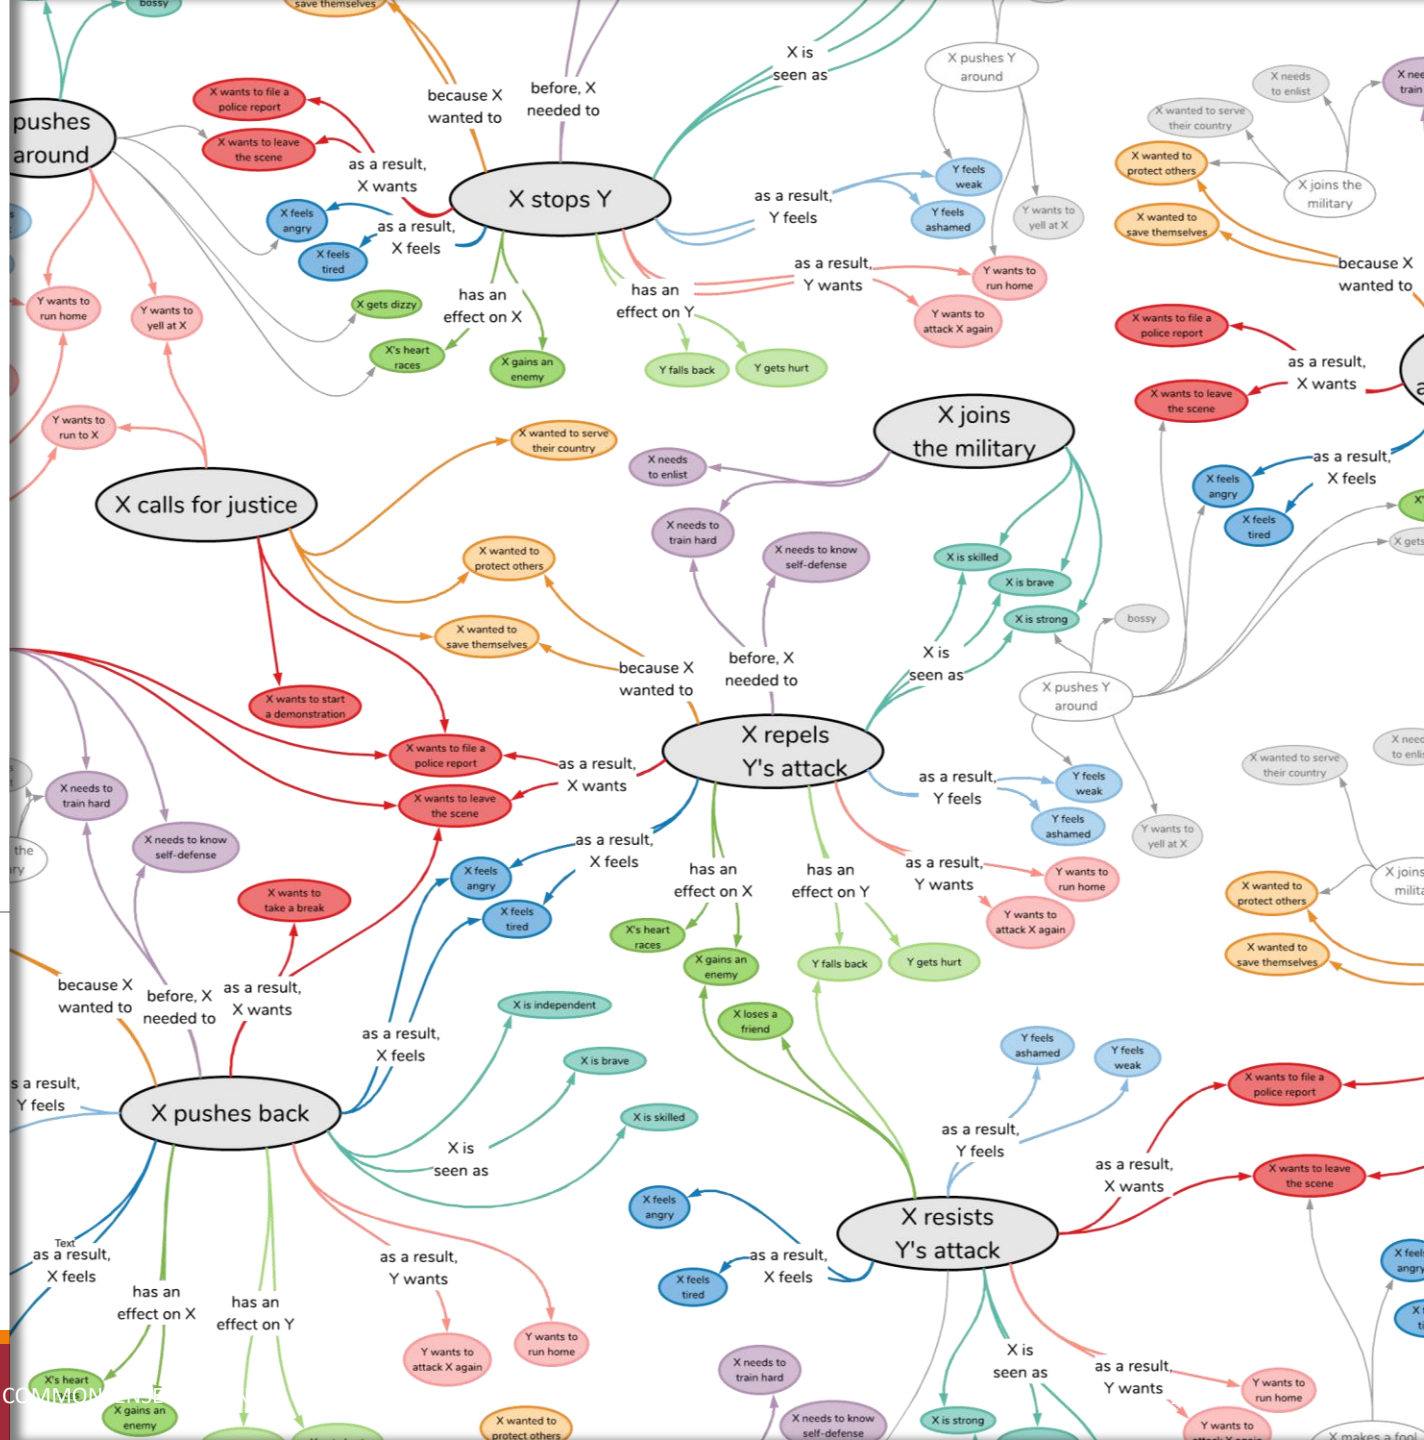
Isaac Asimov's "Three Laws of Robotics"

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey orders given it by human beings except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

- Isaac Asimov, 1942 short story "Runaround"

https://cdn.britannica.com/82/195182-050-97684526/Isaac-Asimov-1979.jpg

Commonsense resources





Grandma’s glasses

Tom’s grandma was reading a new book, when she dropped her glasses.

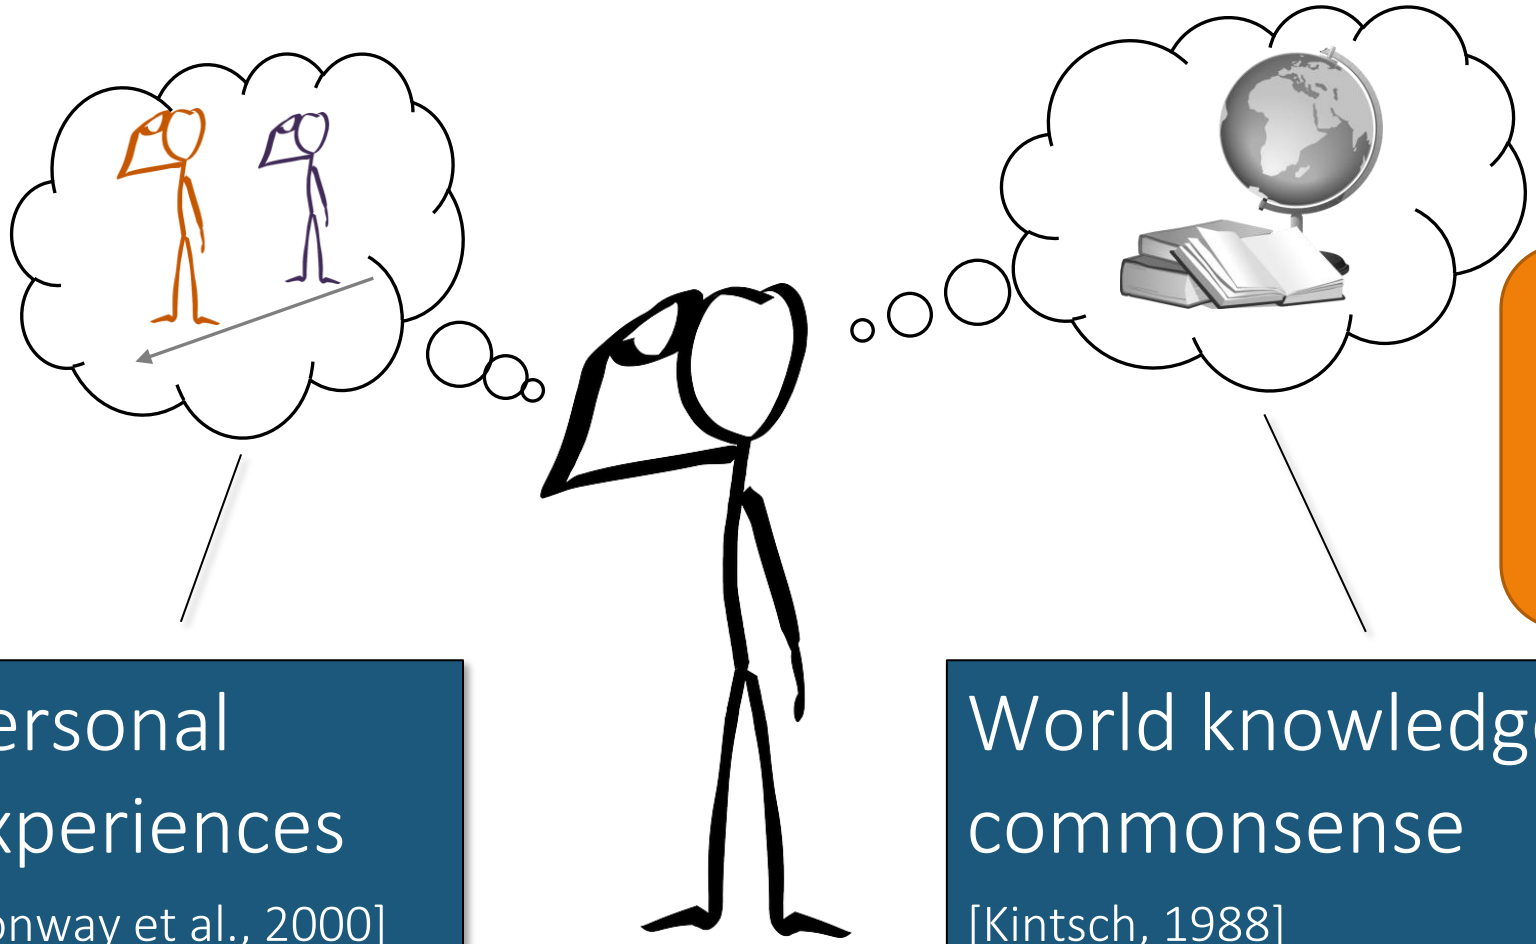
She couldn’t pick them up, so she called Tom for help.

Tom rushed to help her look for them, they heard a loud crack.

They realized that Tom broke her glasses by stepping on them.

Promptly, his grandma yelled at Tom to go get her a new pair.

Humans reason about the world with **mental models**
[Graesser, 1994]



Personal experiences
[Conway et al., 2000]

World knowledge and commonsense
[Kintsch, 1988]

Commonsense resources aim to be a bank of knowledge for machines to be able to reason about the world in tasks

Tom’s grandma was reading a new book, when she dropped her glasses.

She couldn’t pick them up, so she called Tom for help.

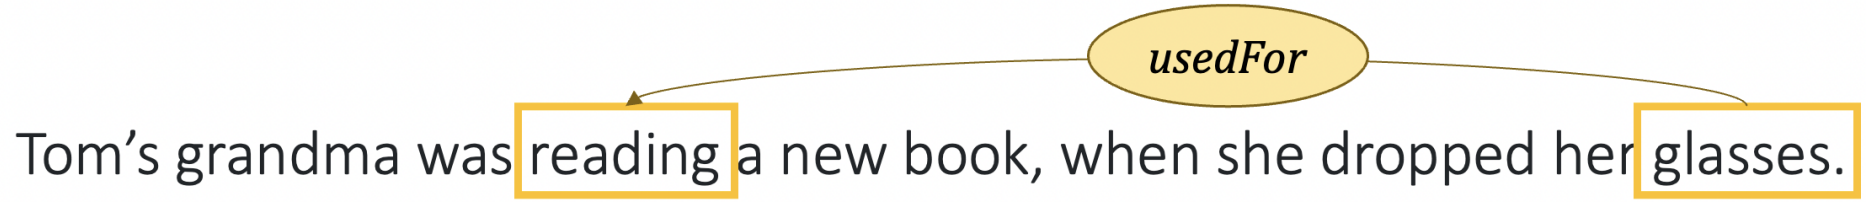
Tom rushed to help her look for them, they heard a loud crack.

They realized that Tom broke her glasses by stepping on them.

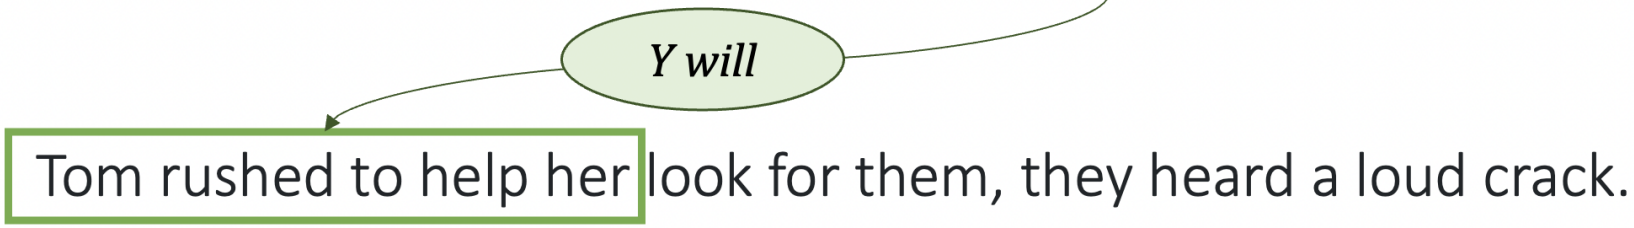
Promptly, his grandma yelled at Tom to go get her a new pair.

ConceptNet

ATOMIC

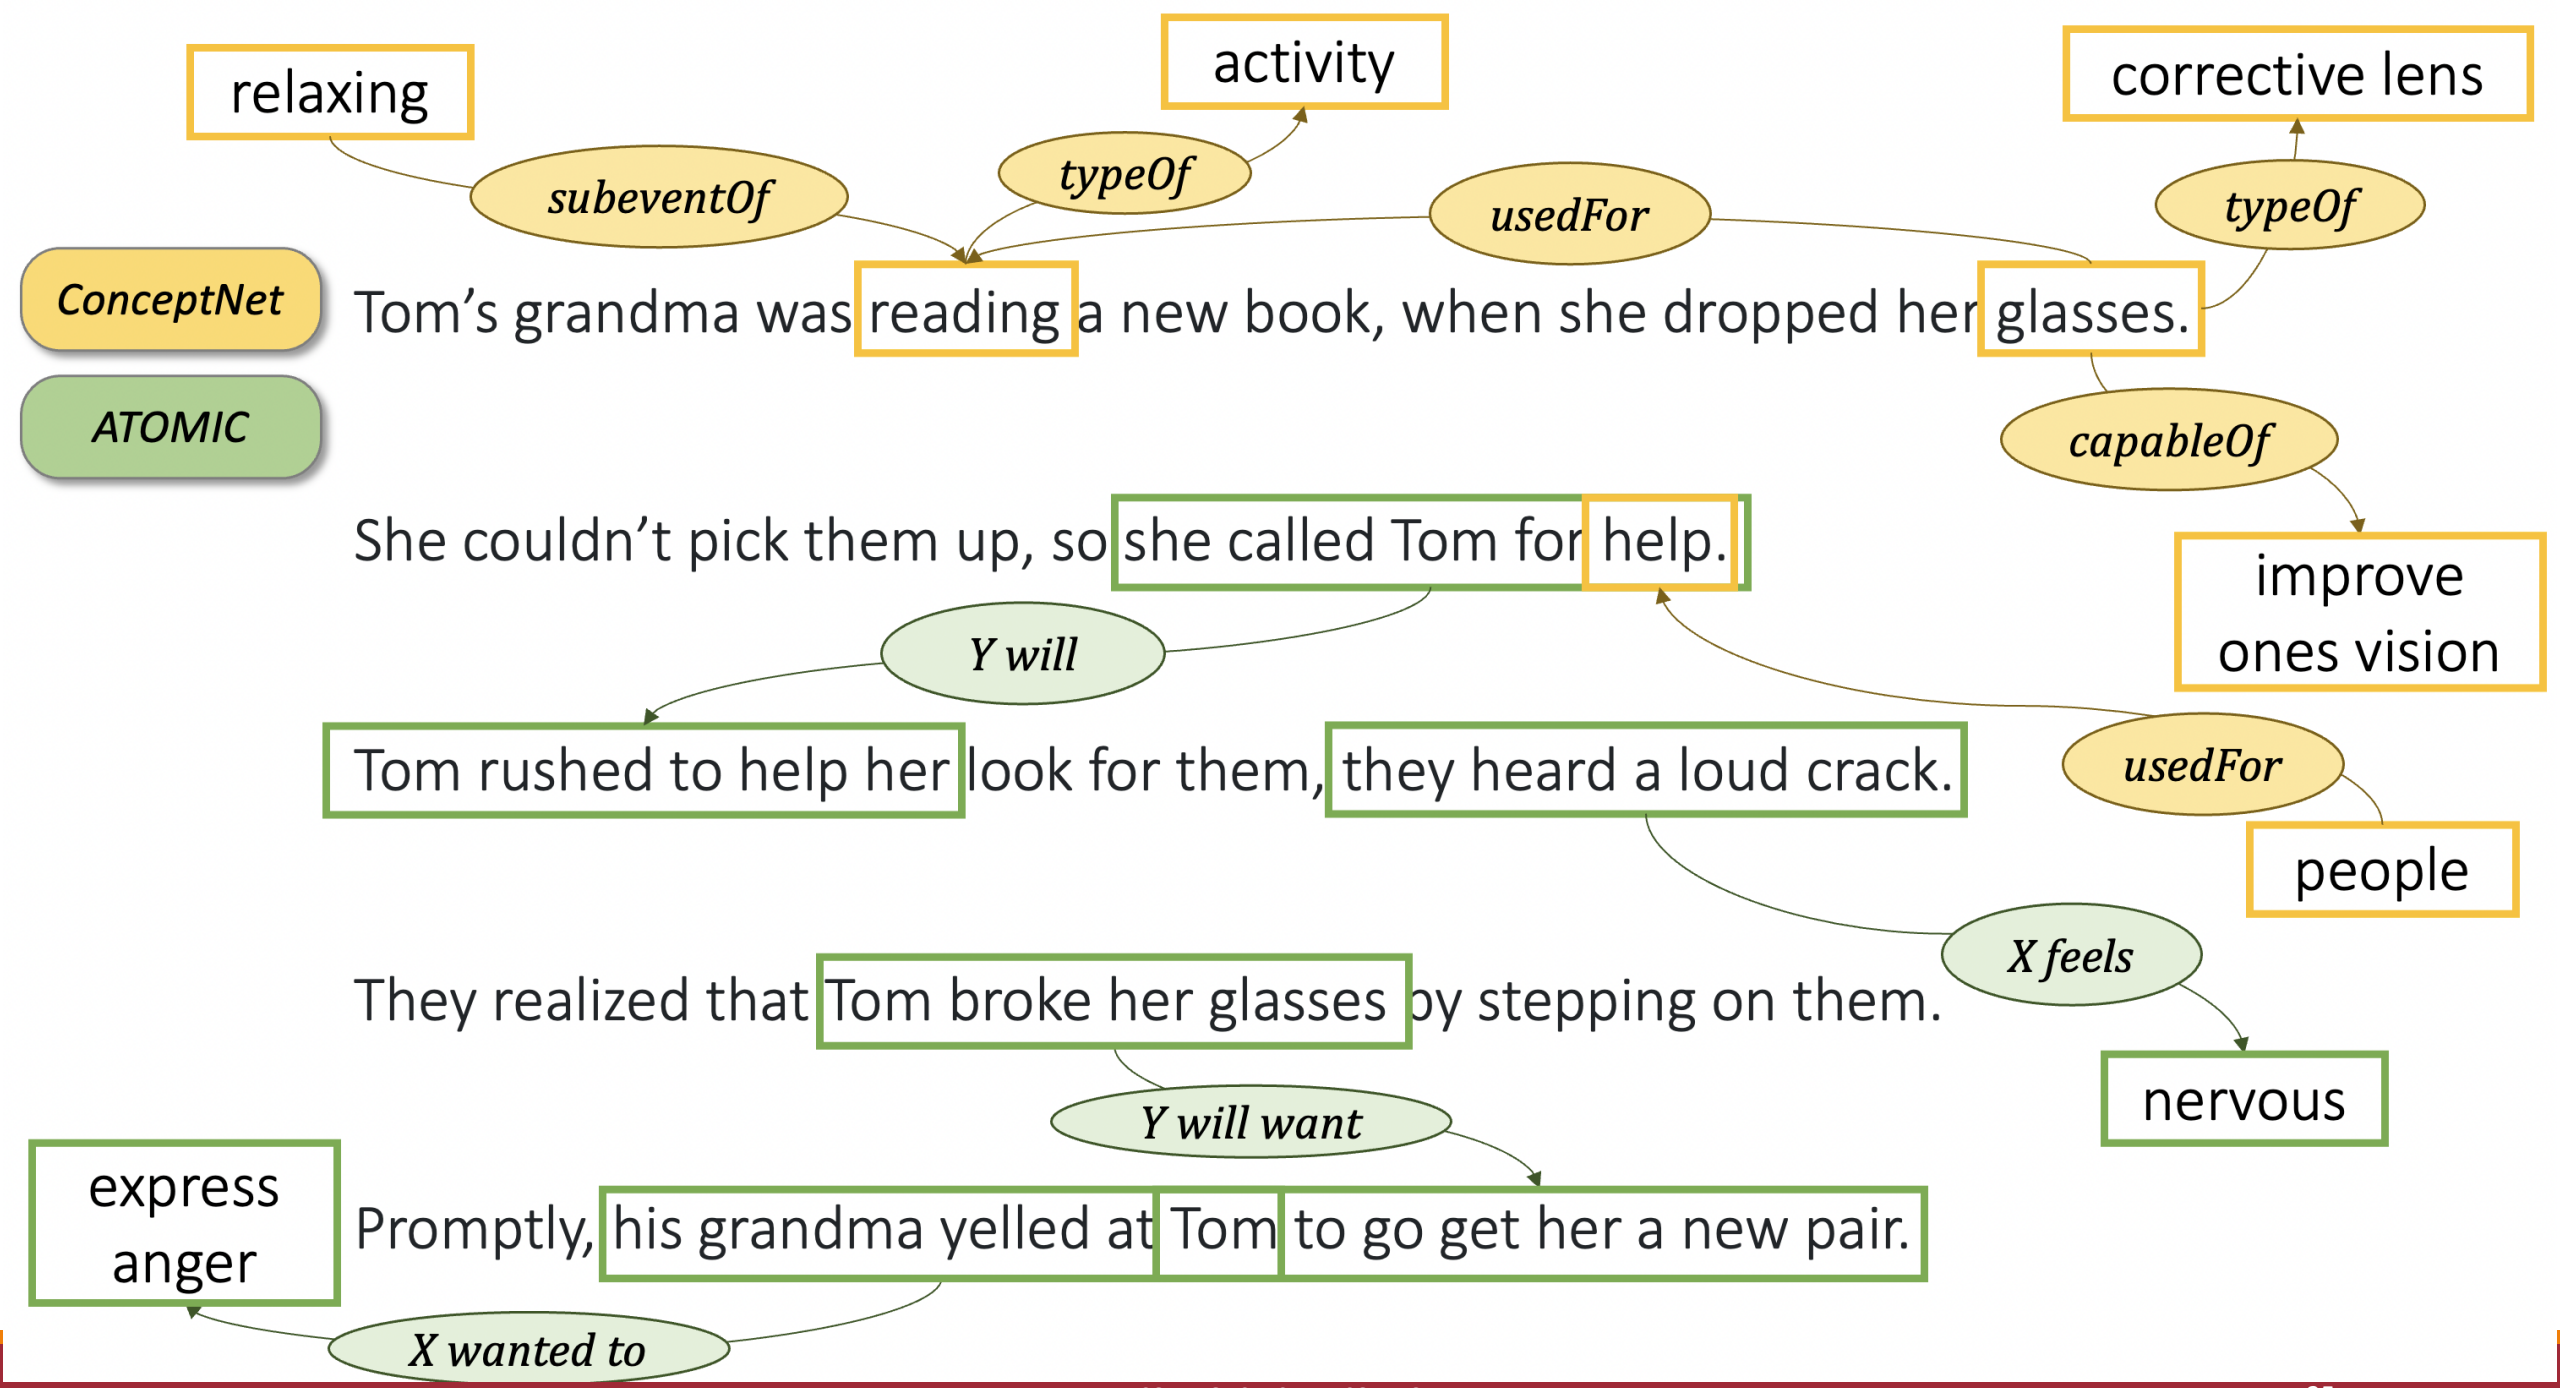


She couldn't pick them up, so she called Tom for help.

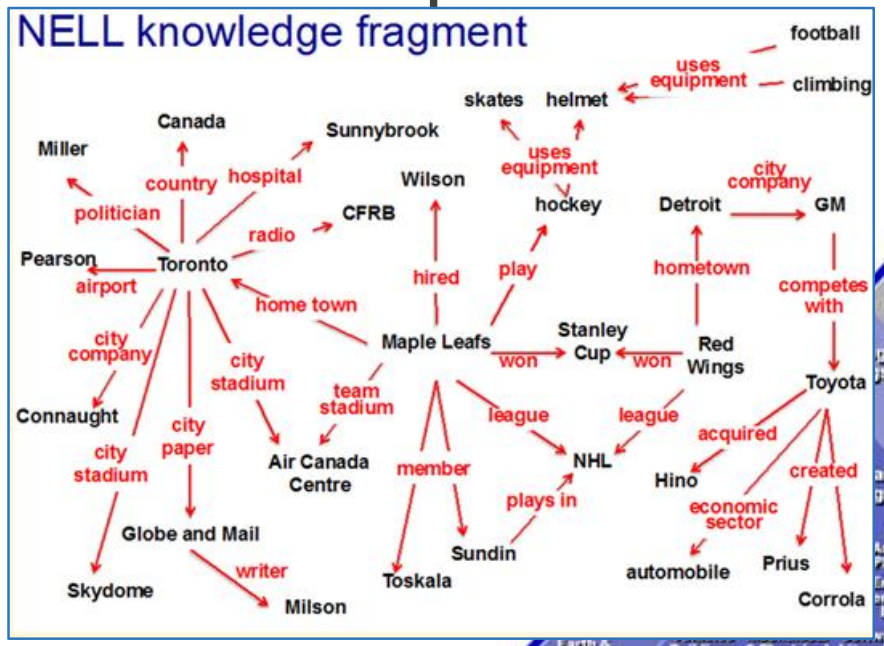


They realized that Tom broke her glasses by stepping on them.




relaxing
activity
corrective lens
subeventOf
typeOf
usedFor
typeOf
ConceptNet
Tom's grandma was reading a new book, when she dropped her glasses.
ATOMIC
capableOf
She couldn't pick them up, so she called Tom for help.
improve ones vision
Y will
Tom rushed to help her look for them, they heard a loud crack.
usedFor
people
X feels
They realized that Tom broke her glasses by stepping on them.
nervous
Y will want
express anger
Promptly, his grandma yelled at Tom to go get her a new pair.
X wanted to

Some important resources



Open Mind Comm
(Minsky, Singh & Ha

ATOMIC
ap et al., 2019)

ild 2.0
al., 2017)

Net 5.5
al., 2017)

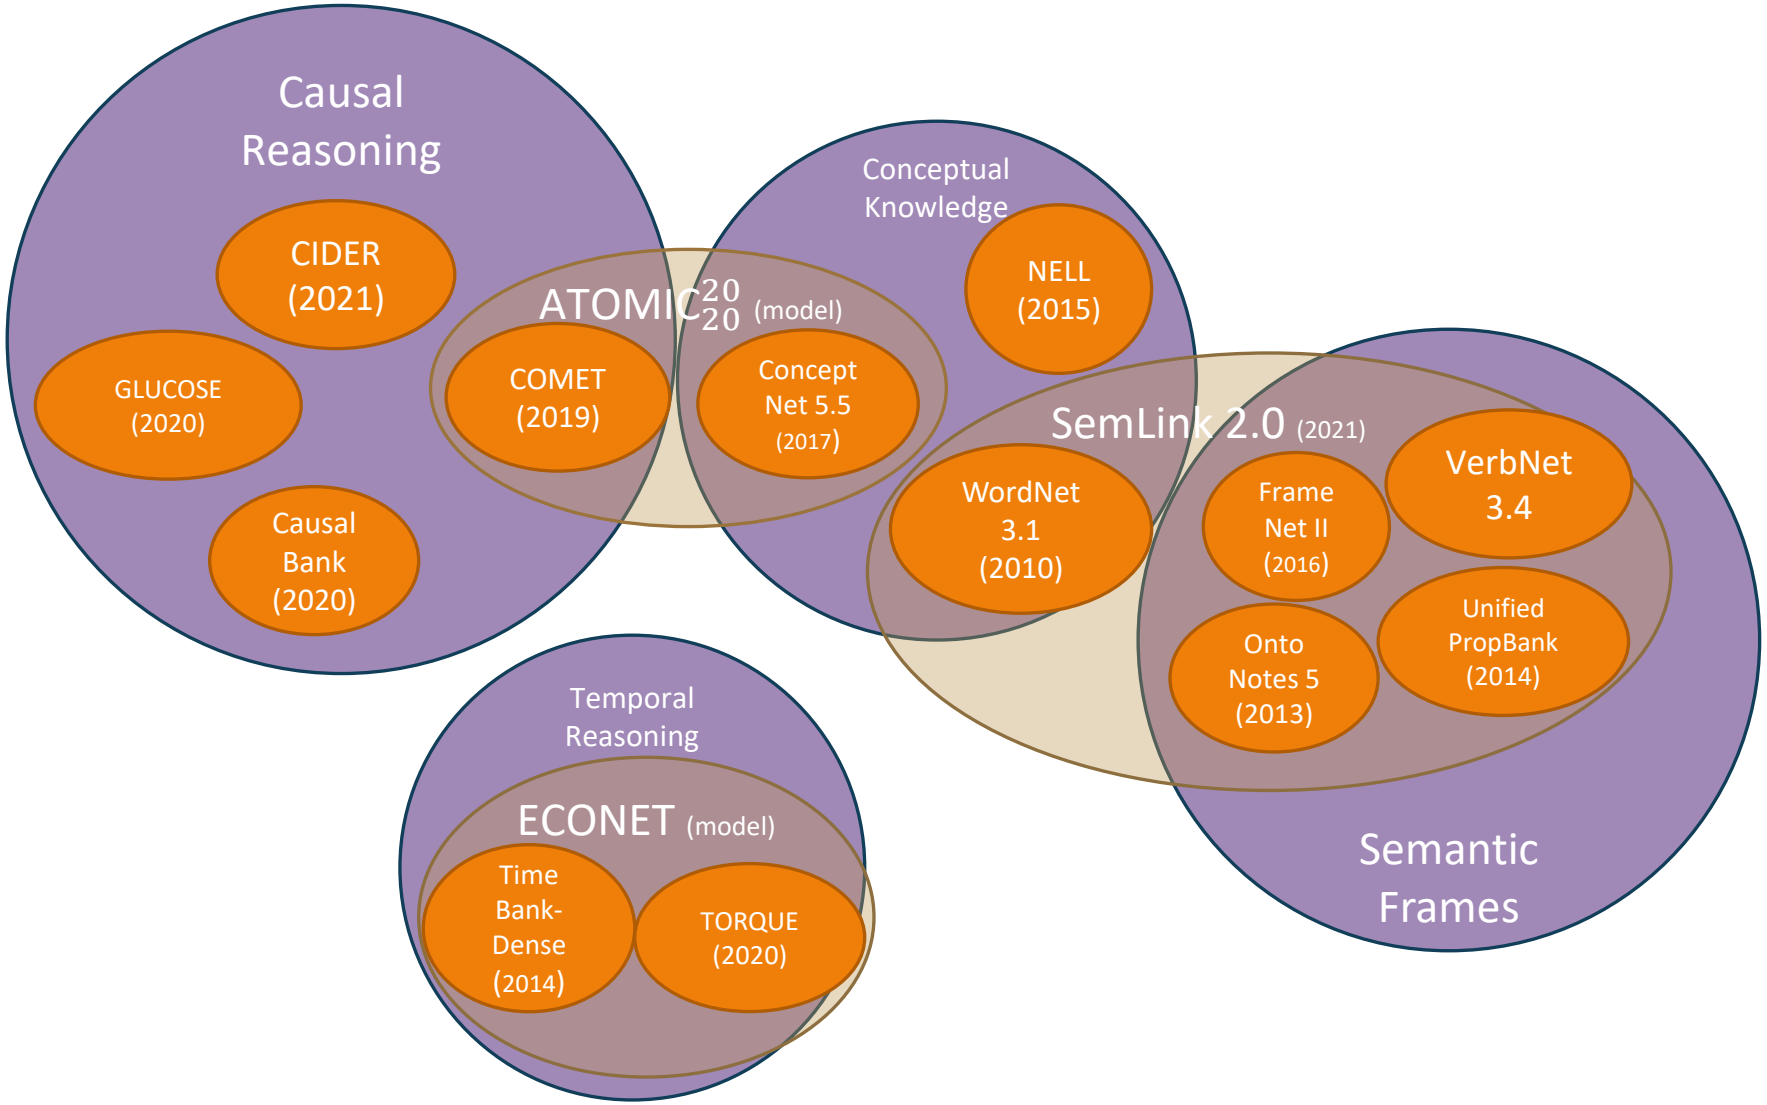
Cyc
(Lenat et al., 1984)

OpenCyc
(Lenat, 2004)

ResearchCyc
(Lenat, 2006)

OpenCyc 4.0
(Lenat, 2012)

today


Causal Reasoning
CIDER (2021)
GLUCOSE (2020)
Causal Bank (2020)
ATOMIC2020 (model)
COMET (2019)
Conceptual Knowledge
Concept Net 5.5 (2017)
NELL (2015)
SemLink 2.0 (2021)
WordNet 3.1 (2010)
Frame Net II (2016)
VerbNet 3.4
Onto Notes 5 (2013)
Unified PropBank (2014)
Semantic Frames
Temporal Reasoning
ECONET (model)
Time Bank-Dense (2014)
TORQUE (2020)