

Commonsense Knowledge Representation and Reasoning in Natural Language processing



Yejin Choi



Vered Schwartz



Maarten Sap



Antoine Bosselut



Dan Roth

Alibaba and Microsoft AI beat human scores on Stanford reading test

Neural networks edged past human scores on the measure of machine reading.

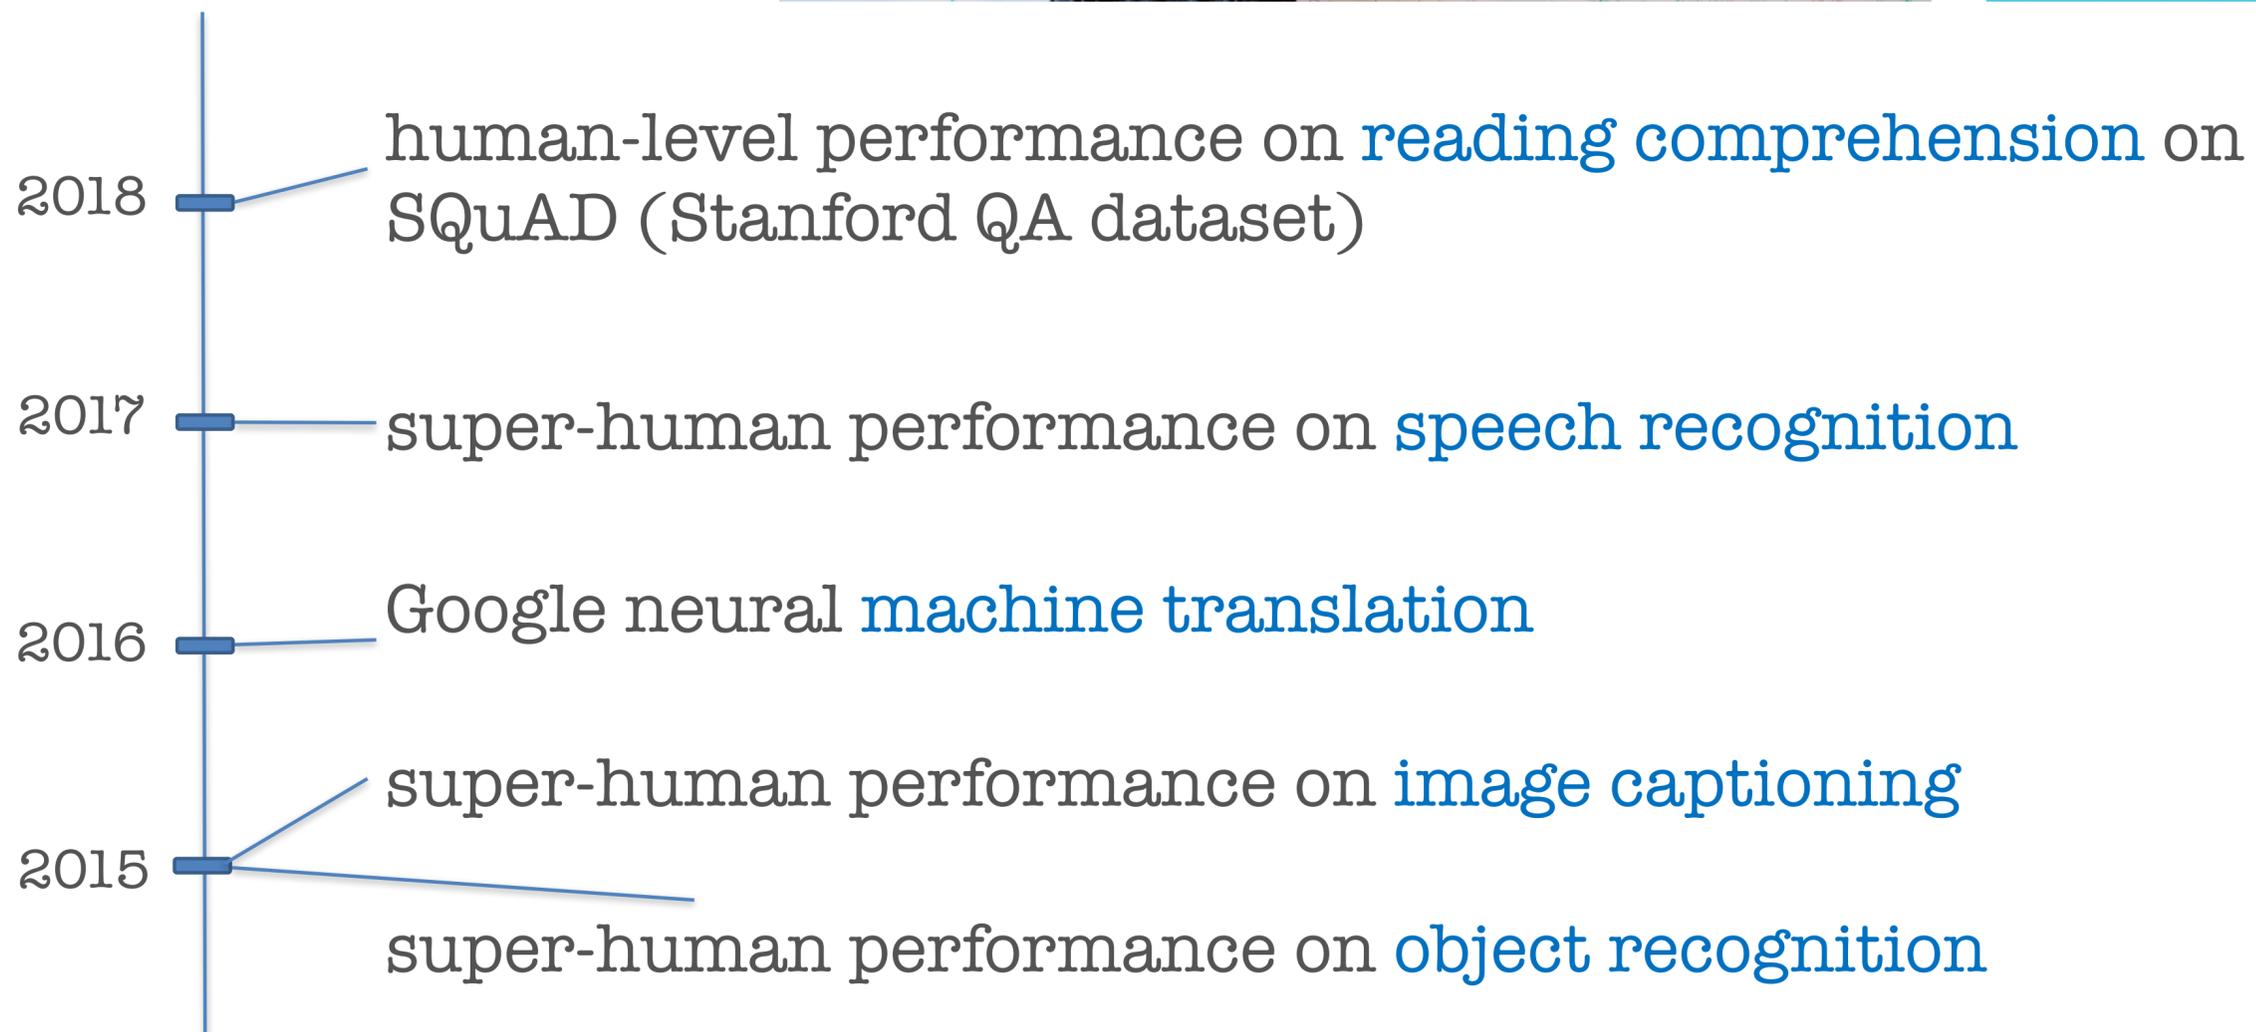
By
f



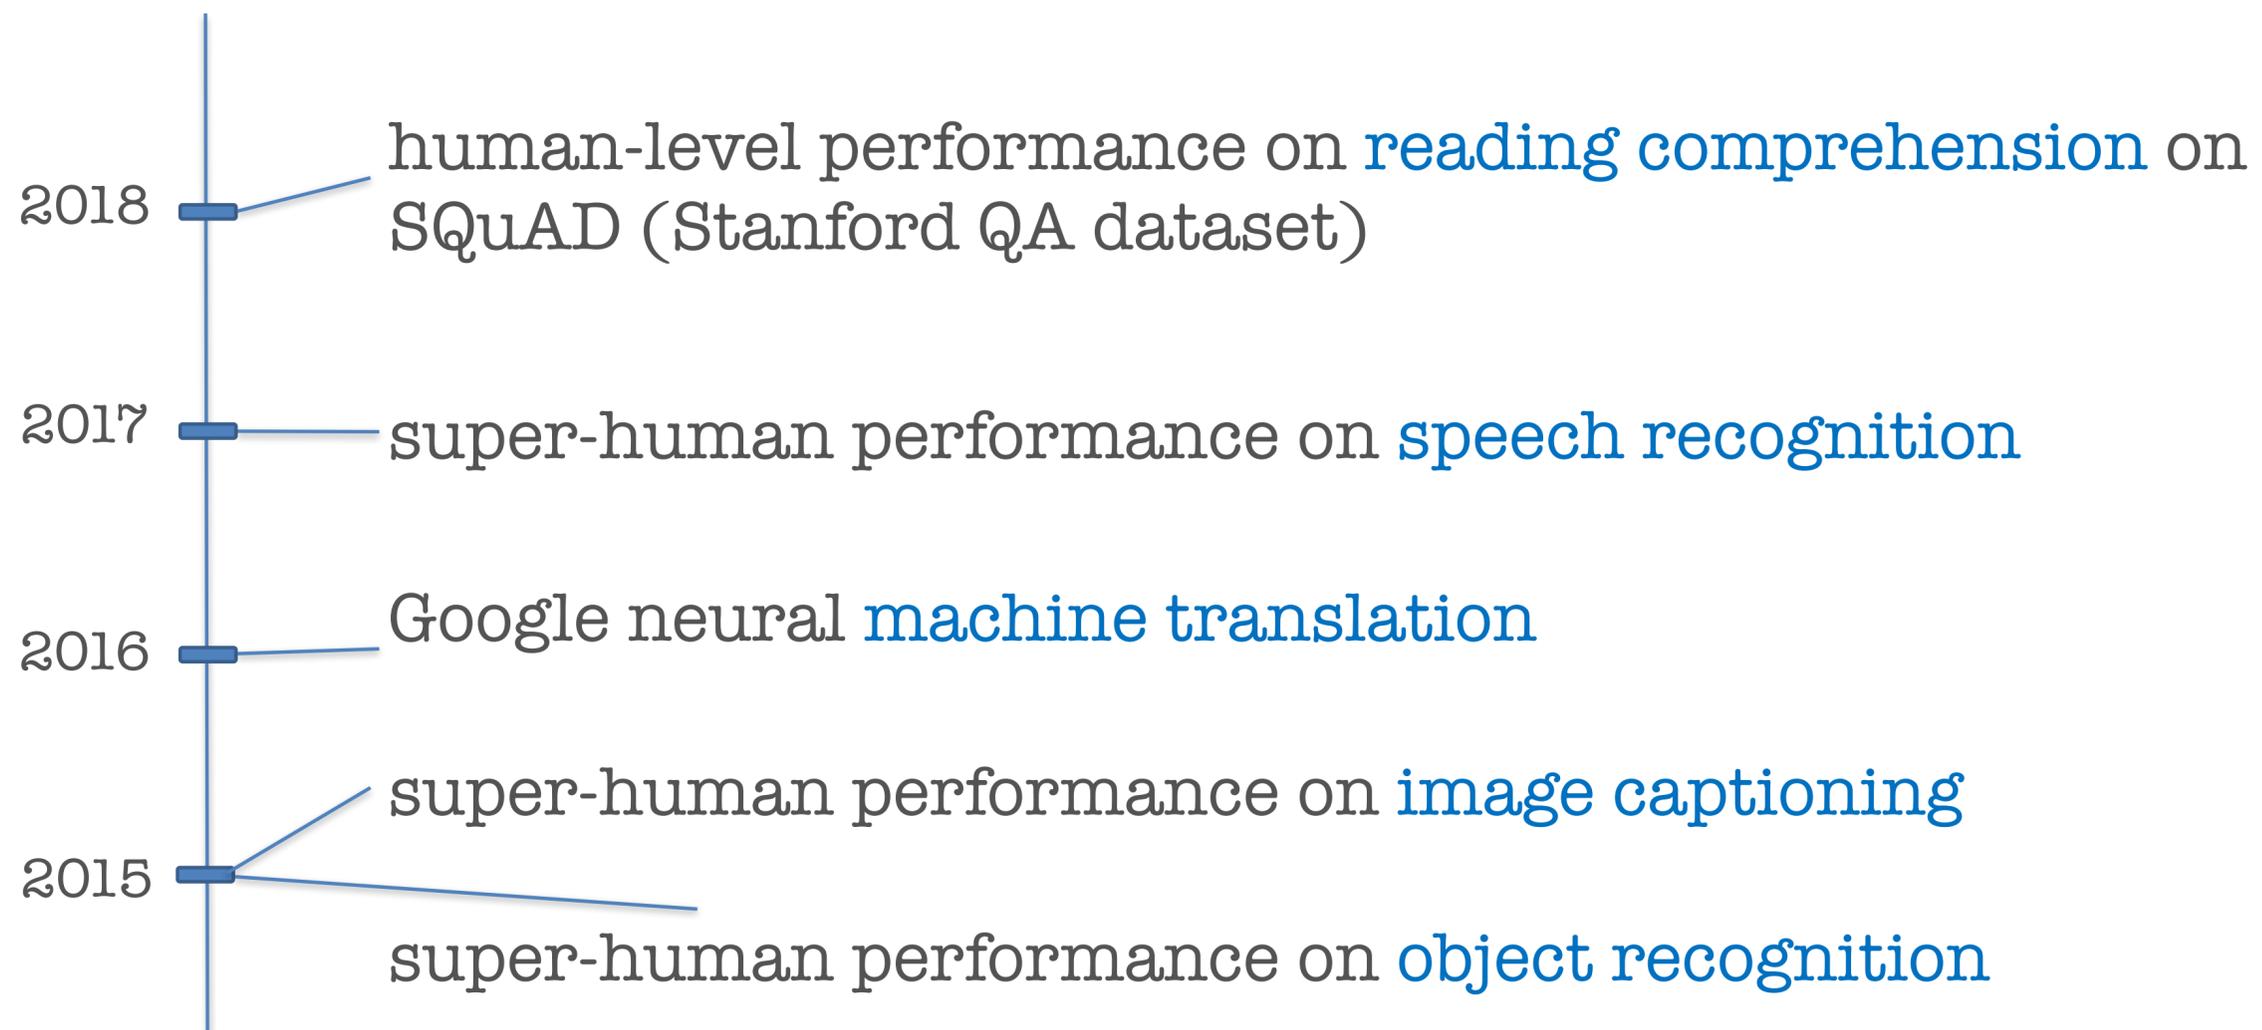
Rob LeFebvre, @roblef
01.15.18 in [Personal Computing](#)

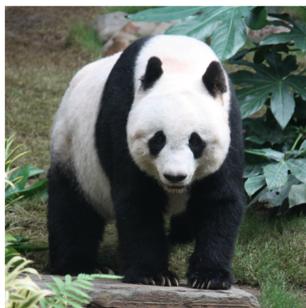
10
Comments

937
Shares

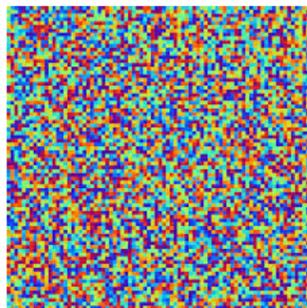


Done solving AI?





+



=



Giant panda
Object

Gibbon

Recognition

Szegedy et al,
2014....



VQA

Jabri et al,
2017

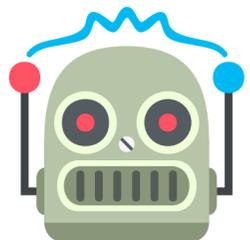


A horse standing in the grass.

Captioning

MacLeod
et al, 2017

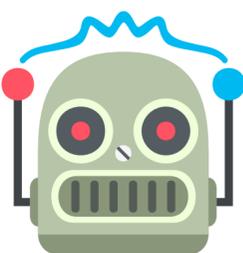
How are you
doing?



I don't know.

Dialogue

Li et al,
2016



I don't know. I
don't know. I
don't know.

Open-ended

Generation

Holtzman
et al, 2018

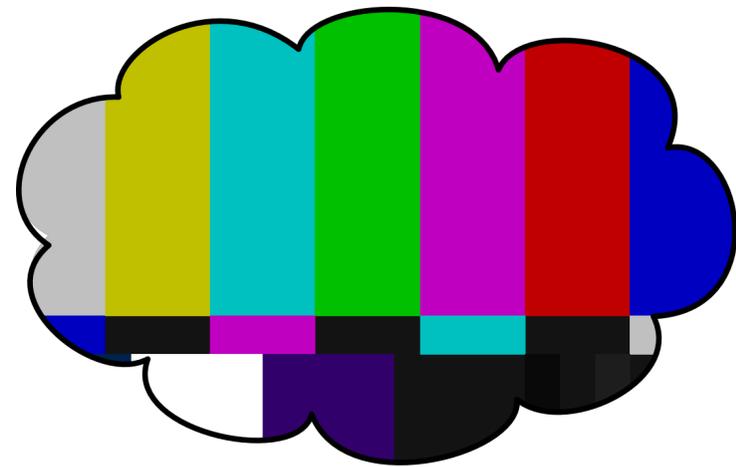
.... Nikola Tesla moved to
Prague in 1880. ... **Tadakatsu**
moved to Chicago in 1881.

Where did Tesla move in
1880? **Chicago**

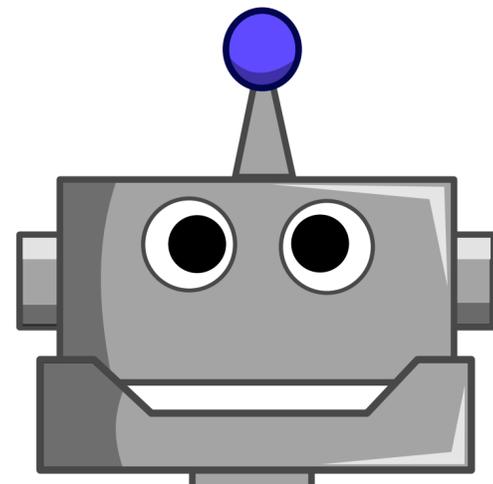
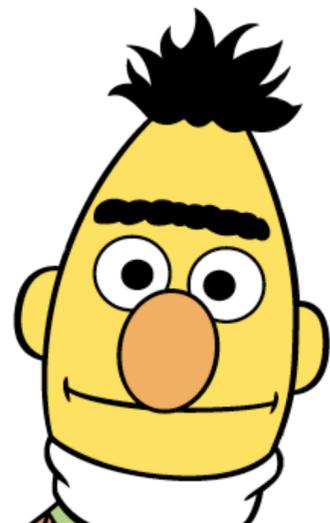
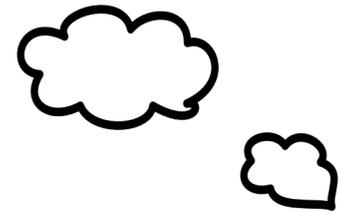
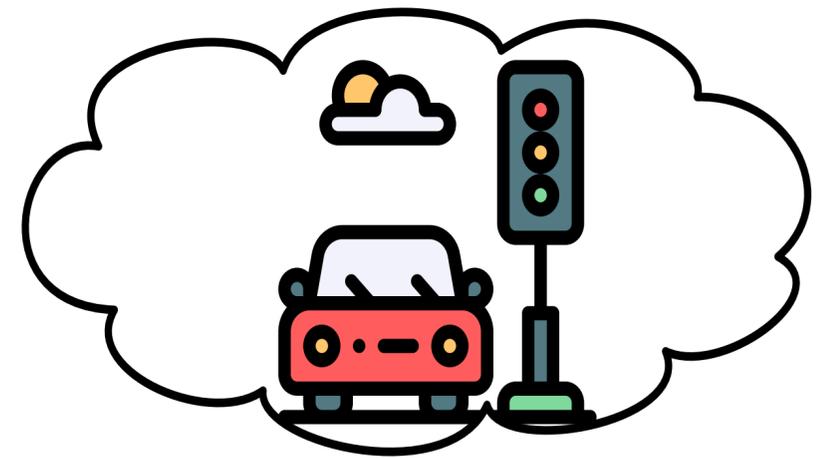
QA

Jia et al,
2017

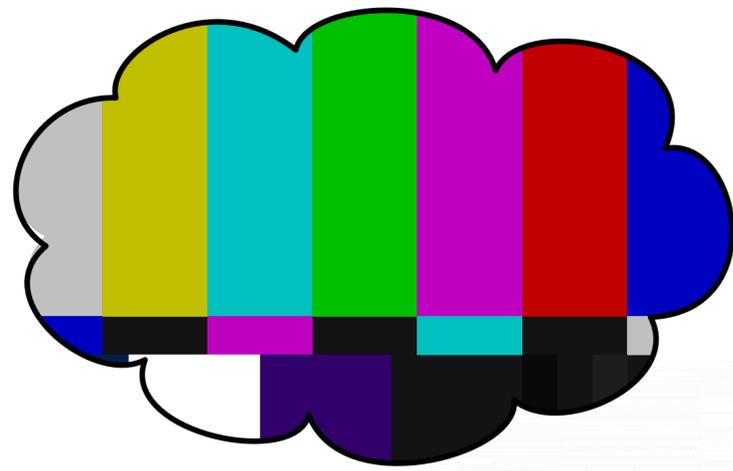
Solving only a "dataset"
without solving the underlying "task"!



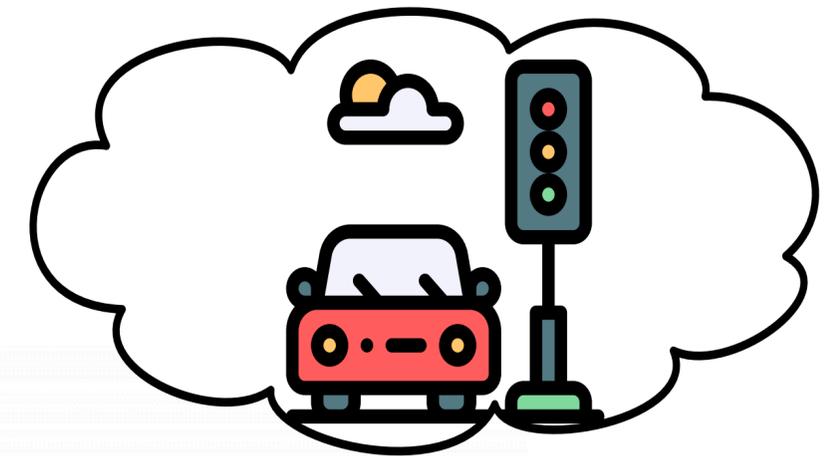
Let's bridge this gap!



*Peters et al., 2018;
Devlin et al., 2018*



Let's bridge this gap!



SYSTEM 1

Intuition & instinct

95%

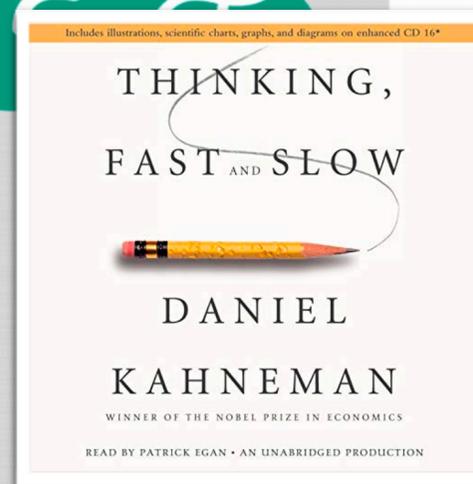
Unconscious
Fast
Associative
Automatic pilot

SYSTEM 2

Rational thinking

5%

Takes effort
Slow
Logical
Lazy
Indecisive

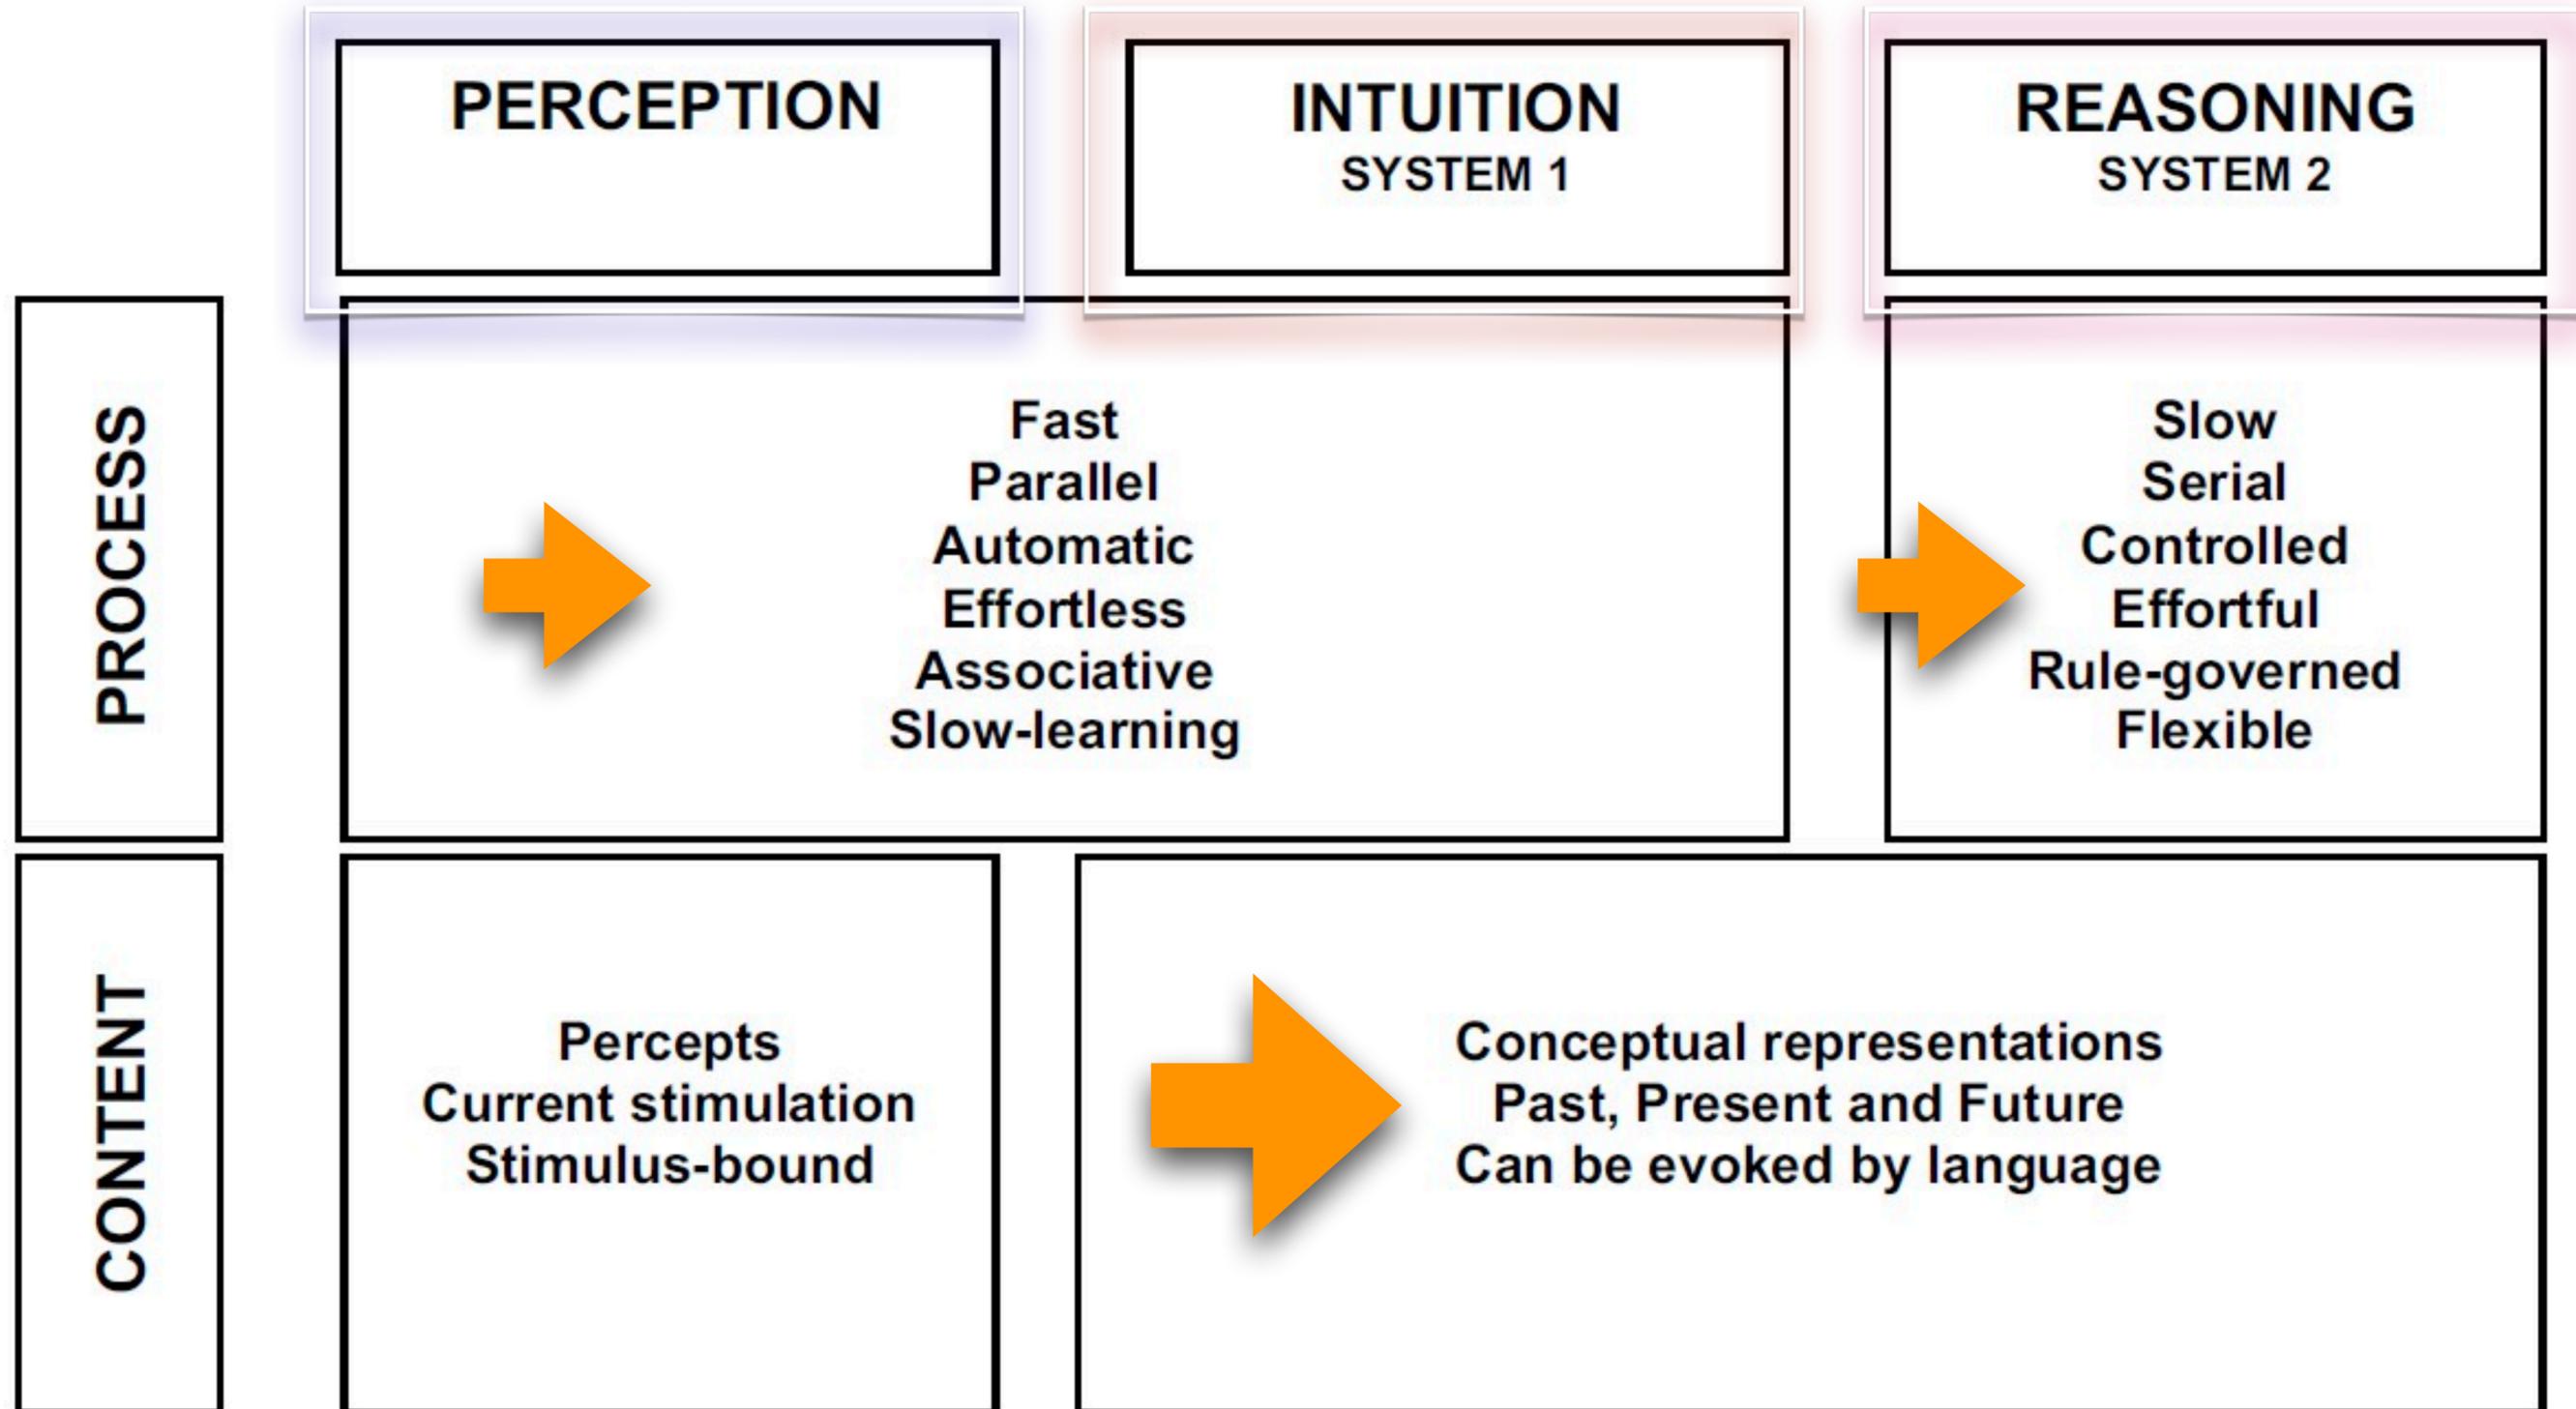


Source: Daniel Kahneman

- **Myth:** we know how to do [system-1 reasoning] with deep learning so we only need to figure out [system-2 reasoning]?

Kahneman's "three cognitive systems"

— "Maps of Bounded Rationality: ..." (Kahneman 2003)



Kahneman's "three **cognitive** systems"

— "Maps of Bounded Rationality: ..." (Kahneman 2003)

PERCEPTION

INTUITION
SYSTEM 1

REASONING
SYSTEM 2

- language models? (predicting one next word ...)
- sequence-to-sequence models? (machine translation, QA, reading comp...)

- object recognition
- image segmentation
- speech recognition

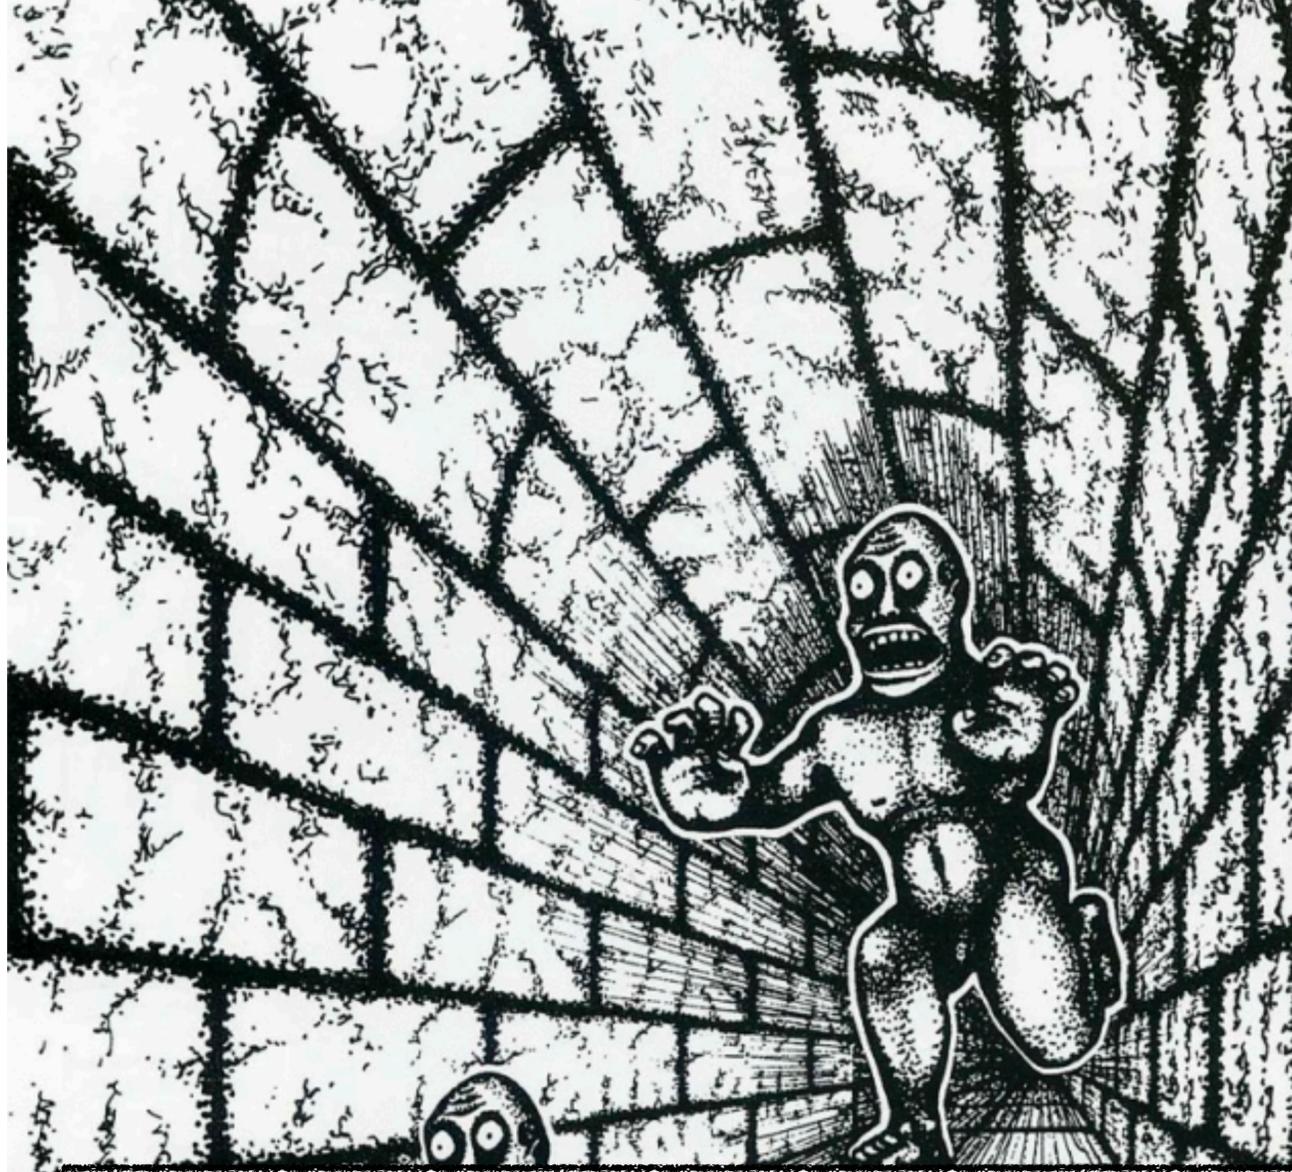
- Intuitive inferences on
 - pre-conditions and post-conditions
 - what happens before and after?
 - motivations and intents
 - mental and emotional states

=> This is what humans do every minute of waking their moments

- solving puzzles
- writing programs
- proving logic theorems

- reviewing ACL papers
- crafting ACL rebuttals
- giving an invited talk
- writing an op-ed

=> Humans often spend hours (or days) **not** doing this sort of reasoning at all...



Roger Shepard's "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)

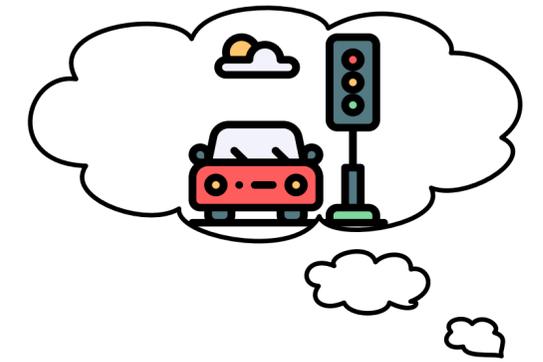
Important Observations:

- 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.

INTUITION
SYSTEM 1

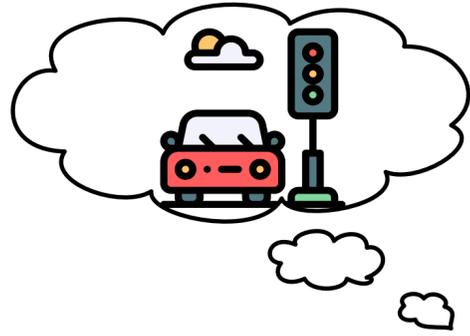
- **Intuitive inferences on**
 - **pre-conditions and post-conditions**
 - **what happens before and after?**
 - **motivations and intents**
 - **mental and emotional states**

Definition of Common Sense

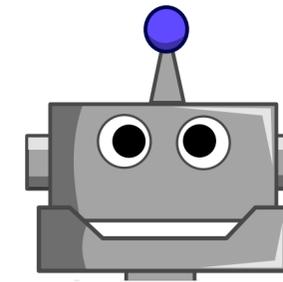
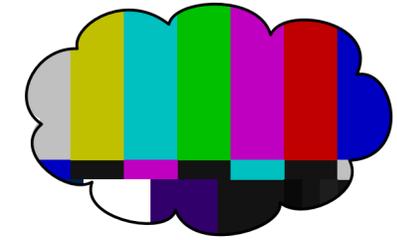


- the basic level of **practical knowledge** and **reasoning** 🤪
- concerning **everyday situations** and **events**
- that are **commonly** shared among **most** people.

For example, it's ok to keep the closet door open,
but it's not ok to keep the fridge door open,
as the food inside 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

For example, it's ok to keep the closet door open, but it's not ok to keep the fridge door open, as the food inside might go bad.

Commonsense

- Searching “commonsense” from ACL anthology
 - Most papers are either from 80s or from the past few years

Position Paper on Common-sense and Formal Semantics

Geoffrey Nunberg
Xerox PARC and CSLI, Stanford

1. A philological excursus

I'm not sure what I'm doing on this panel, but I thought it would be helpful if we could start at the beginning. It's interesting to note that both the dictionary and common sense were eighteenth-century inventions. This is no coincidence; in fact, it's entirely appropriate that the most celebrated

Revisiting Commonsense

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

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

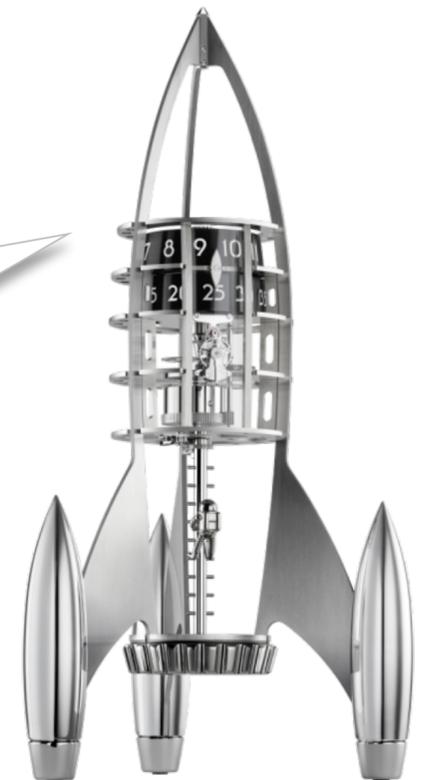
- weak computing power
- not much data
- no crowdsourcing
- not as strong computational models
- not ideal conceptualization / representations

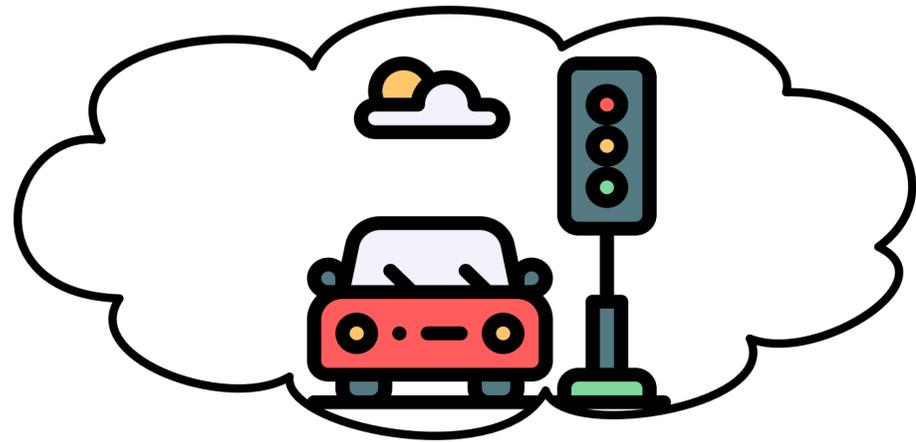
Path to commonsense?



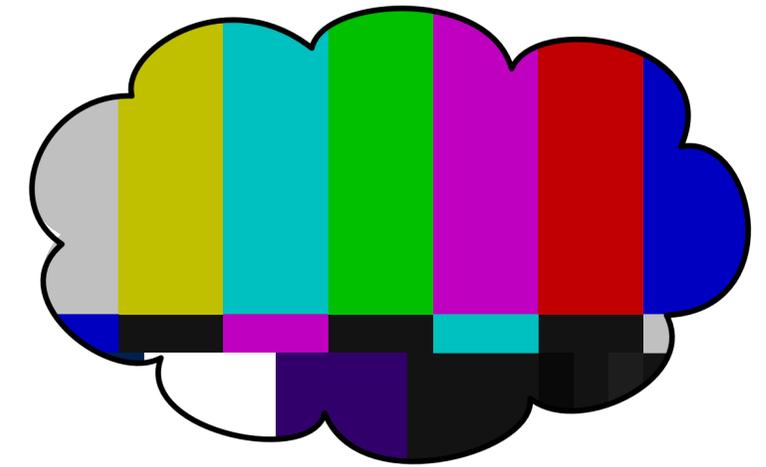
Brute force larger networks with deeper layers?

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





Let's bridge this gap!



Symbolic
common
sense
graph

Neural
commonsense
representations

Reasoning
engine with
common
sense

Constructing
challenge
datasets
right



Tutorial Schedule

Approximate Time (PT)	Segment Title	Speaker	Description
15:00 - 15:15	Introduction	Yejin Choi	Why commonsense reasoning is the new frontier of artificial intelligence
15:15 - 15:40	Knowledge in LMs	Vered Shwartz	On the types of commonsense knowledge captured during the pre-training of language models, and what is still missing
15:40 - 16:10	Commonsense Resources	Maarten Sap	How to gather and represent commonsense knowledge of different types (e.g., social, physical, taxonomic)
16:10 - 16:30	Commonsense Integration into Neural Networks Part 1	Vered Shwartz	How to enhance neural models for commonsense reasoning tasks with symbolic knowledge
	Break		
17:00 - 17:30	Commonsense Integration into Neural Networks Part 2	Antoine Bosselut	How language models can be converted to commonsense knowledge bases, and the downstream effects of these new tools
17:30 - 17:55	Commonsense Benchmarks	Maarten Sap	How to create benchmarks to measure whether and how well a model can do commonsense reasoning
17:55 - 18:25	Temporal Common sense	Dan Roth	Studies in discovering the temporal commonsense implications of events described in text
18:25 - 18:30	Concluding Remarks	Dan Roth	Thoughts and Future Directions

Introduction



"Commonsense reasoning is the new frontier of artificial intelligence. Let's chat about these implications."

15:00 – 15:15 PT

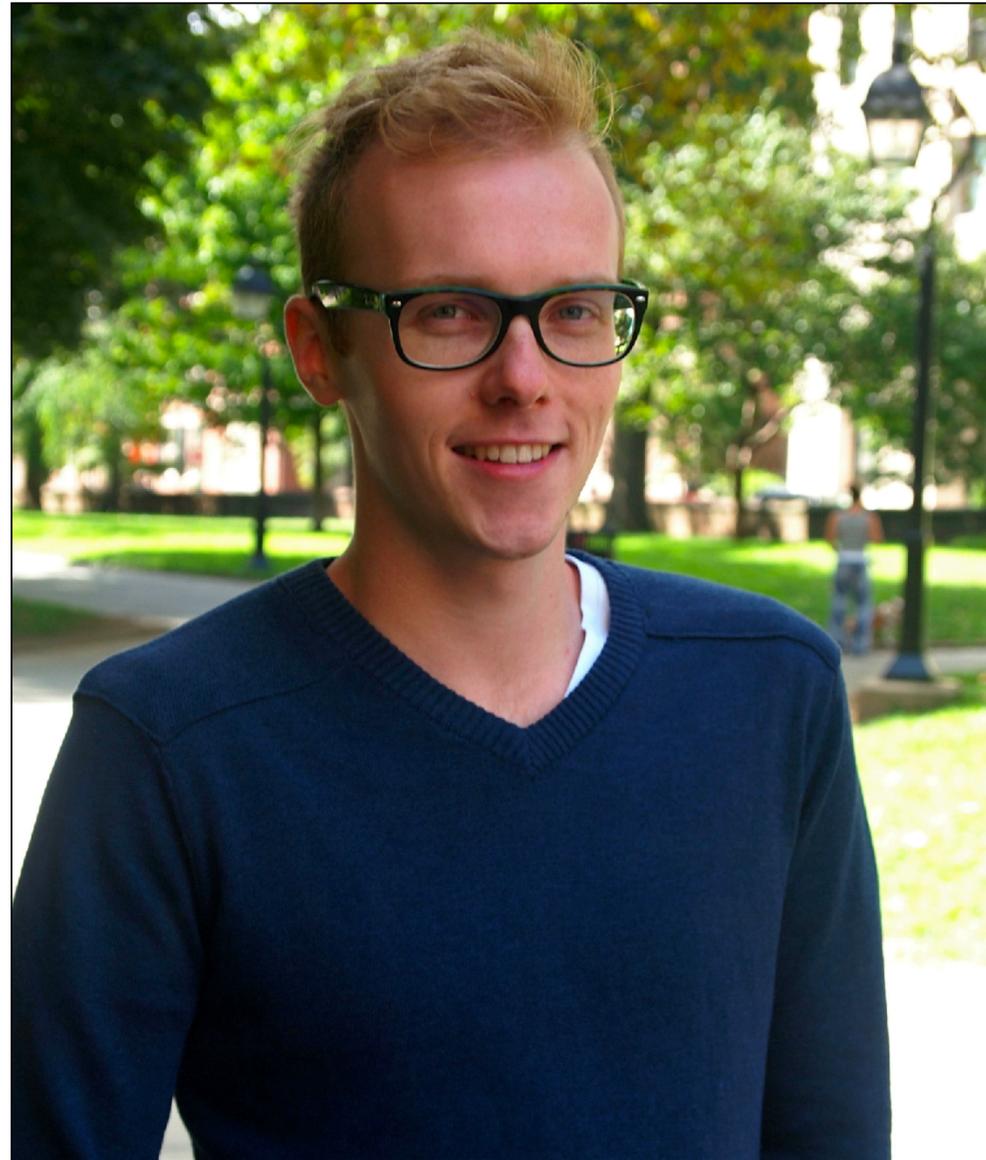
Knowledge in Language Models



“Let’s look at the types of commonsense knowledge that are captured and missing in pretrained language models.”

15:15 – 15:40 PT

Commonsense Resources



“Let’s learn to gather and represent commonsense knowledge of different types – social, physical, taxonomic.”

15:40 – 16:10 PT

Integrating Commonsense into Neural Networks



“Let’s enhance neural models for commonsense reasoning tasks with symbolic commonsense knowledge.”

16:10 – 16:30 PT

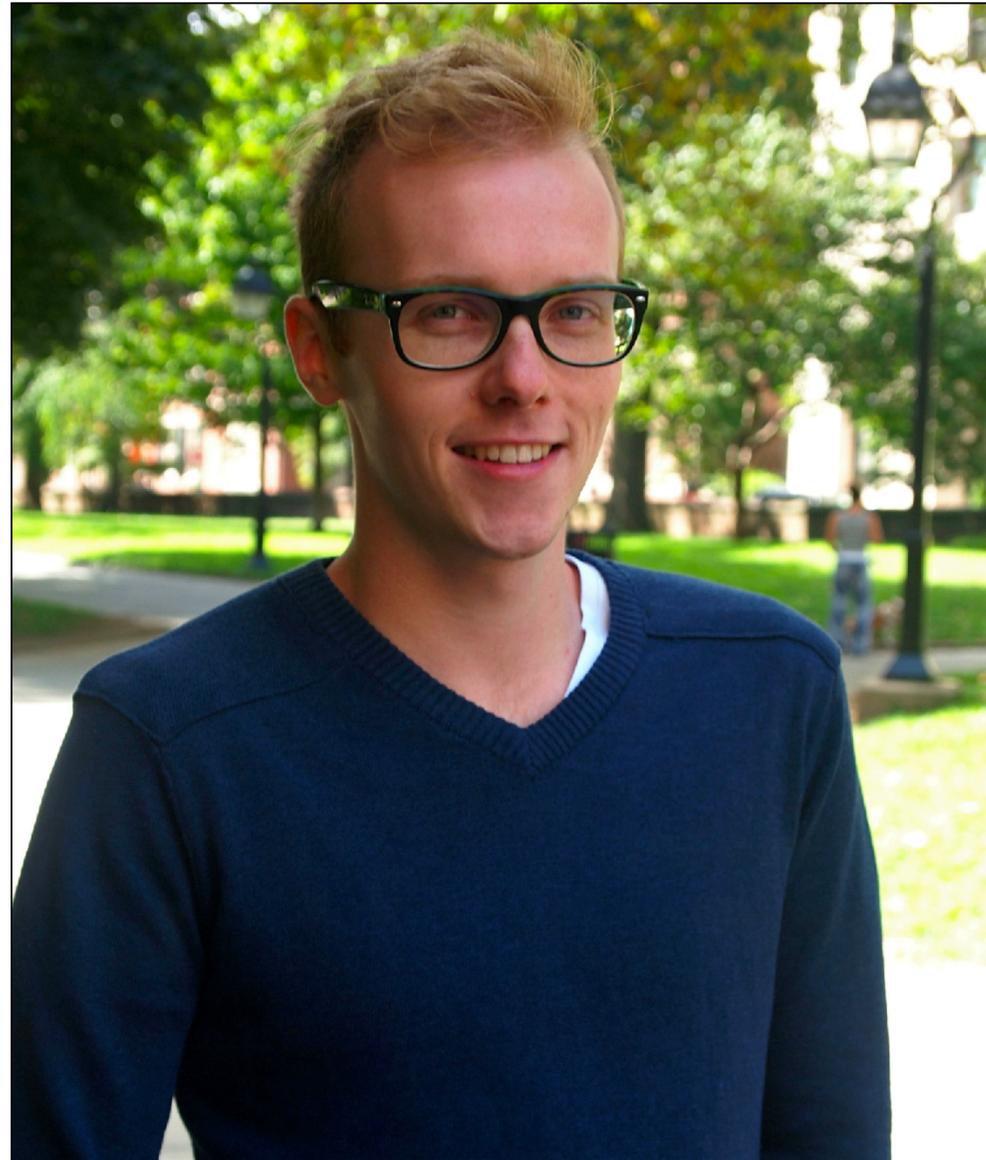
Integrating Commonsense into Neural Networks



“Let’s use language models as commonsense knowledge bases, and look at the effect on NLP pipelines.”

17:00 – 17:30

Commonsense Resources



“Let’s create new benchmarks to measure how well models can reason about commonsense knowledge.”

17:30 – 17:55 PT

Temporal Commonsense Reasoning



“Let’s talk about the importance of understanding commonsense implications of time.”

17:55 – 18:25 PT

Conclusion



“Let’s talk about next steps, and how commonsense representation will change NLP research.”

18:25 – 18:30 PT