

Exploring Transitivity in Neural NLI Models through Veridicality

Hitomi Yanaka^{1,2} Koji Mineshima³ Kentaro Inui^{4,1}

¹RIKEN, ²University of Tokyo, ³Keio University, ⁴Tohoku University



Abstract

Question: To what extent neural models can learn the systematicity in language from training instances?

- We analyze the systematic generalization ability of neural models in the domain of Natural Language Inference (NLI) on **transitivity inference** with synthetic and naturalistic datasets involving **veridicality**

Transitivity

Key concept for Systematicity [Fodor and Pylyshyn, 1988] in NLI

- If a model learns $A \rightarrow B$ and $B \rightarrow C$, it should compose the two and draw a new inference $A \rightarrow C$

Otherwise, a model must memorize every possible combination independently!

Dataset creation

- **Veridical inference** can easily compose transitivity inference at scale by embedding various inferences into clause-embedding verbs
- We select 30 veridical (e.g. *s knows P* entails *P*) / non-veridical (e.g. *s thinks P* does not entail *P*) verbs *f* from previous veridicality datasets

[White+2018; Ross and Pavlick, 2019] and embed synthetic/naturalistic inference $s_1 \rightarrow s_2$ into *f*

Synthetic datasets

veridical + synthesized Boolean inference (*and, or, not*)

$f(s_1)$: *Someone knew that [Bob found Tom, Jim and Fred]*

s_1 : *Bob found Tom, Jim and Fred*

s_2 : *Bob found Jim*

Naturalistic datasets

veridical + lexical/structural inference (SICK [Marelli+, 2014])

$f(s_1)$: *Someone thinks that [a person is brushing a cat]*

s_1 : *A person is brushing a cat*

s_2 : *A person is combing the fur of a cat*

yes : Entailment
unk : Non-entailment

$f(s_1) \rightarrow s_1$	$s_1 \rightarrow s_2$	$f(s_1) \rightarrow s_2$
yes	yes	yes
yes	unk	unk
unk	yes	unk
unk	unk	unk

Evaluation protocol

Training

Basic 1. veridical inference: $f(s_1) \rightarrow s_1$

Premise: *Jo {knew/hoped} that Bob and Ann left. [f(s1)]*

Hypothesis: *Bob and Ann left. [s1]* (Entailment/Non-entailment)

Basic 2. various inference patterns: $s_1 \rightarrow s_2$

Premise: *Bob and Ann left. [s1]*

Hypothesis: *Ann left. [s2]* (Entailment)

Test transitivity inference involving veridicality: $f(s_1) \rightarrow s_2$

Premise: *Jo {knew/hoped} that Bob and Ann left. [f(s1)]*

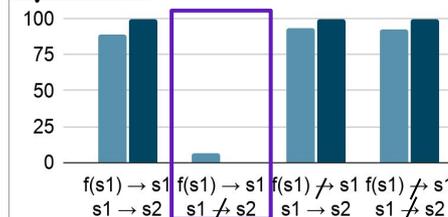
Hypothesis: *Ann left. [s2]* (Entailment/Non-entailment)

Conclusion

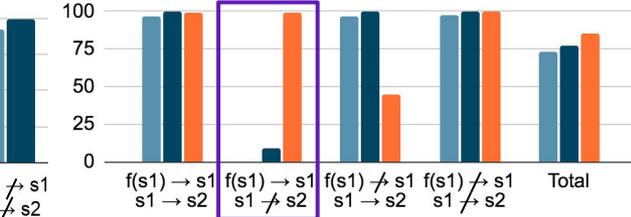
Current NLI models fail to consistently perform transitivity inference, lacking the generalization capacity for drawing composite inferences from provided training examples. There remains room for improving the systematic generalization capacities of models with respect to combining basic inferential abilities.

Main results

Synthetic



Naturalistic



- The models do not perform well on the cases where the verb *f* is veridical but embedded sentence s_1 does not entail s_2 .
⇒ only look at the veridicality of *f*.
- Humans tend to fail on the cases where *f* is non-veridical and s_1 entails s_2 .
⇒ only look at the relation between s_1 and s_2 , neglecting the veridical verb. (Veridicality bias [Ross and Pavlick, 2019])

Contact: hitomi.yanaka@riken.jp
<https://github.com/verypluming/transitivity>