

TTIC 31230, Fundamentals of Deep Learning

David McAllester, Autumn 2020

AGI: Natural Language

Natural Language for Knowledge Representation

Is natural language a reflection of a general learning or knowledge representation architecture?

Will some Transformer-like architecture ultimately underly general intelligence?

Natural Language for Knowledge Representation

I like to sample news stories to try understand how language represents reality.

Example: President Donald Trump on Tuesday insisted he was serious when he revealed that he had directed his administration to slow coronavirus testing in the United States, shattering the defenses of senior White House aides who argued Trump's remarks were made in jest.

Natural Language for Knowledge Representation

I take the position that reference, rather than compositional meaning, is the fundamental semantic phenomenon in language.

Most of the phrases of a sentence refer to a particular entity or event.

We can introduce constant symbols for the entities and events and break sentences down into simple statements involving the entities.

Natural Language for Knowledge Representation

President Donald Trump on Tuesday insisted he was serious when he revealed that he had directed his administration to slow coronavirus testing in the United States, shattering the defenses of senior White House aides who argued Trump's remarks were made in jest.

[the order] = [Trump] directed his administration to reduce [covid testing].

[the revelation] = [Trump] revealed [the order].

[the claim] = [the order] was in jest.

[the assertion] = [the aides] made [the claim].

Natural Language for Knowledge Representation

President Donald Trump on Tuesday insisted he was serious when he revealed that he had directed his administration to slow coronavirus testing in the United States, shattering the defenses of senior White House aides who argued Trump's remarks were made in jest.

[the aides] = Senior White House aides

[the defense] = [the assertion] defended [Trump].

[the insistence] = on Tuesday [Trump] insisted [the order] was serious.

[the shattering] = [the insistence] shattered [the defense].

Subtle Semantics

We would like to be able to represent the subtle semantics of words like “insisted”, “serious”, “revealed”, “shattered”, “defenses” and “argued”.

A translation to first order logic would introduce predicate symbols for these words.

However, their subtle meanings seem impossible to express in the rigid Boolean semantics of first order logic.

Neural Semantics

In contrast to Boolean semantics, the meaning of words seems to require embeddings processed by neural networks.

But as in logic, discrete “statements” involving words and entities seem important for representing reality.

[the insistence] = on Tuesday [Trump] insisted [the order] was serious.

Natural Language in Code and Mathemamtics

Mnemonic names and natural language comments are essential in writing intelligible code or mathematical proofs.

This supports the idea that natural language somehow underlies our understanding of what would seem to be purely logical constructs.

Rather than logic providing the meaning of language, neural language processing may support our understanding of logic.

Or maybe logic and neural language processing just support each other.

END