NS4NLP: Closing Remarks

Maria L. Pacheco, Sean Welleck, Yejin Choi, Vivek Srikumar, Dan Goldwasser, Dan Roth

COLING 2022 Tutorial

Oct. 2022

Neuro (and/vs.) Symbolic Approaches

- **The Neuro-Symbolic distinction is often characterized as**
  - **System 1:** thinking fast*, typically associated with learned neural models*
  - **System 2:** thinking slow*, typically associated with symbolic reasoning*

- Large focus on **perceptual tasks** (System I):
  - Consuming lots of raw data, assign categories to them

- Many real-world scenarios have **unique challenges**
  - Need to understand relationships between higher-level variables
  - Fairness, Accountability, Transparency

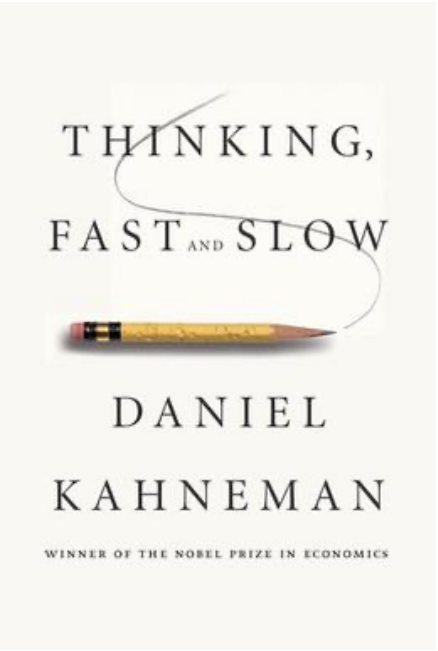

Ongoing Debates in the NLP Community

- NAACL 2022 Panel: ***The place of linguistics and symbolic structures***

- **We cannot deny the capabilities of Neural Language Models,** but their limitations are unclear

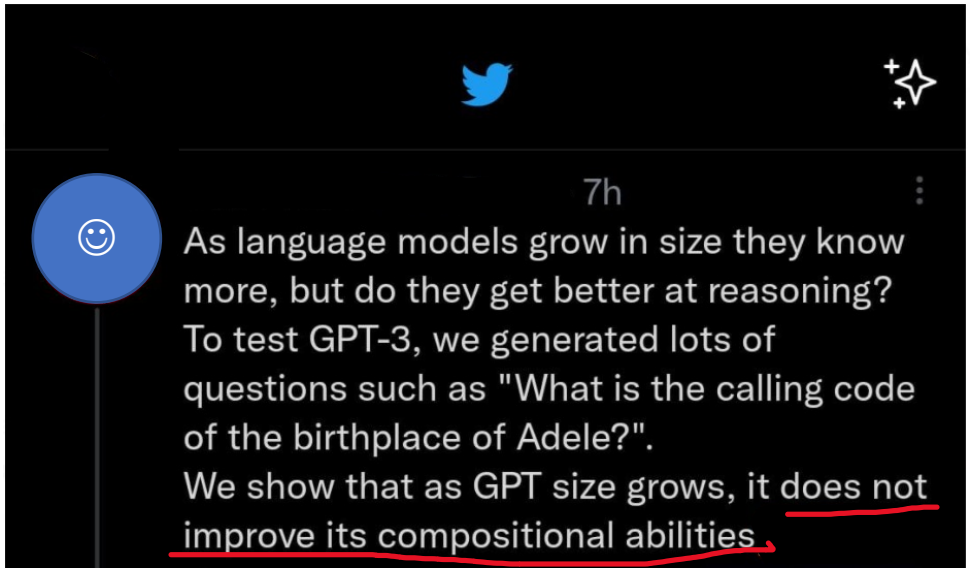

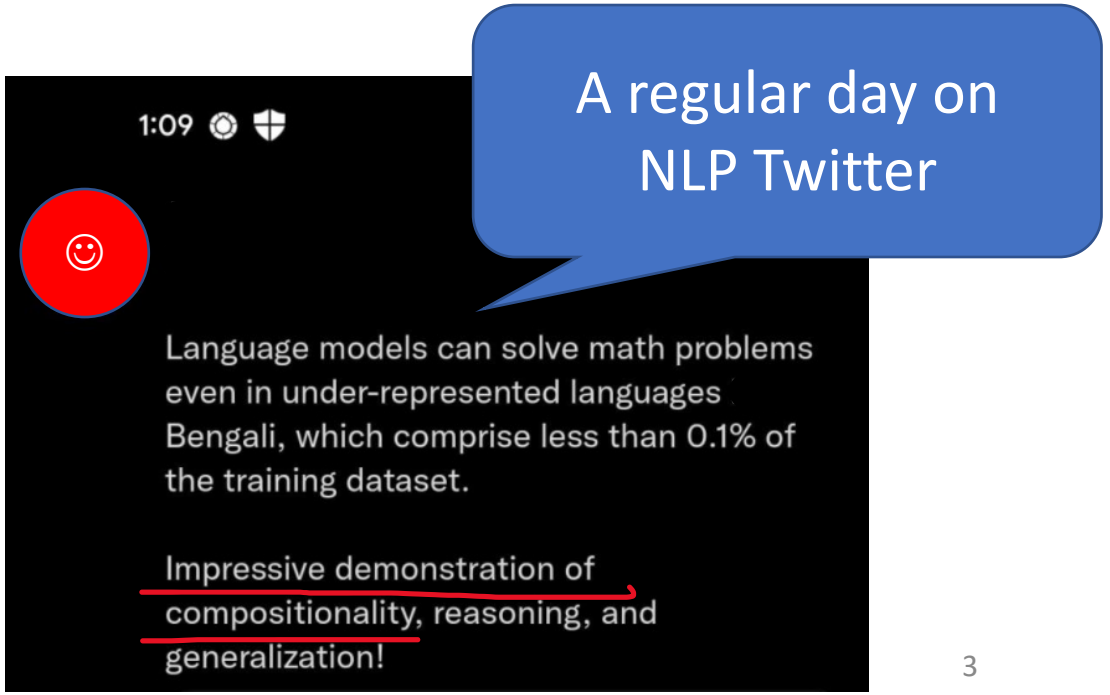

Challenges of Purely Neural Models

- **It is more-or-less accepted that purely neural models face challenges**
  - Energy efficiency
    - *LLMs are prohibitive for most.*
  - Data efficiency
    - *And the unintended consequences of swallowing the internet.*
  - Explainability
    - *If we do not understand them, how can we trust them?*
  - Human Interaction
    - *Making use of the knowledge that we have.*
  - Reasoning beyond surface level patterns
    - *We want to avoid picking up spurious correlations*

Challenges turned Opportunities

- **Neuro-Symbolic methods:** ***Are they the best of both worlds?***

Neural methods can train expressive models using massive datasets to identify patterns in raw data

Symbolic methods can map neural representations to symbols and reason over higher level patterns.
They are inherently explainable and easier to manipulate.

What is Next for NS4NLP?

- Hopefully we did a good job of convincing you to care, **but what is next?**

- **There are still many problems to tackle!**

  - What is the "*right*" neuro-symbolic representation? Is there one?

  - What is the right level of abstraction? What can be deferred to the NNet?

  - Scaling and performance – combinatorial inference is expensive, but so is backprop!

  - The role of symbols - an interface for Human-AI interaction or abstractions for internal reasoning?

Thank You!

Oct. 2022

Demo

Analyzing Political Discourse

- Tweets written by US congress members on the ***abortion*** issue
- Use morality frames (moral foundations, roles)

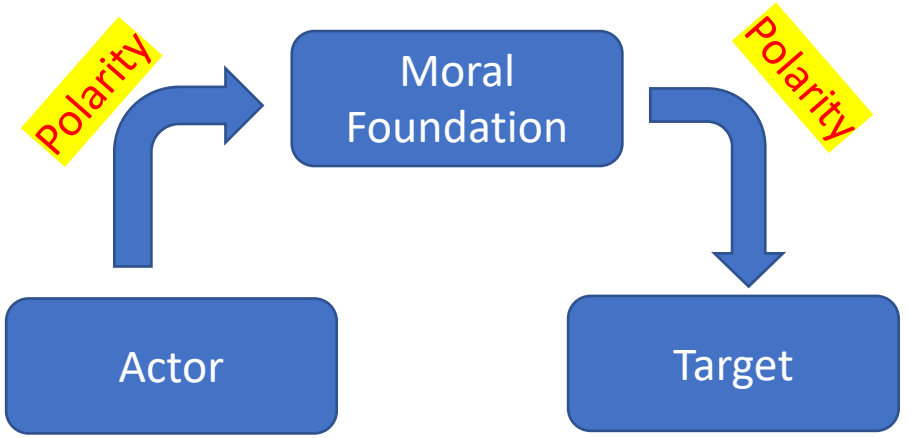

- Opinions range from “pro-life above all”, to “women’s choices above all”
- We will look at the ***entities*** at the center of this debate
  - Women, babies, life, US government institutions, legislative bills

Morality Frames to Analyze The Abortion Debate

- ***Entities***

Event -> *One Global Instance*

Tweet

Topic -> *Abortion*

Ideology -> *Left, Right*

Entity-Mention -> *"The ACA"*

Entity-Group -> *ACA*

Moral Foundation -> *Fairness, Care, etc.*

Role -> *target-care*, *provide-care, etc.*

Polarity -> *positive, negative*

- ***Relations***

InInstance(Tweet, Event)

HasEntity(Tweet, Entity)

HasTopic(Tweet, Topic)

HasIdeology(Tweet, Ideology)

HasRole(Entity, Role)

RoleHasMF(Role, MF)

RoleHasPolarity(Role, Polarity)

Using DRaiL to Analyze the Abortion Debate

Base Classifiers:

- InInstance(T, Z) & HasEntity(T, E) => HasRole(T, E, R^RoleLabel?)
- InInstance(T, Z) => HasMf(T, M^MfLabel?)

Party Messaging Preferences:

- InInstance(T, Z) & HasIdeology(T, I) => HasMf(T, M^MfLabel?)
- InInstance(T, Z) & HasEntity(T, E) & HasIdeology(T, I) => HasRole(T, E, R^RoleLabel?)

Joint Inference:

- InInstance(T, Z) & HasEntity(T, E) & RoleHasMf(R, M) & HasRole(T, E, R)^? => HasMf(T, M)^?

Live Demo