

Linguistics Lab

13 March 2024 • NLU

On NLP/ML/AI/etc and Linguistics

Lukewarm take: NLP and Linguistics are not the most compatible fields

= linguistics + ML people ask different questions, don't work well together

Chapter 1

Modern language models refute Chomsky's approach to language

Steven T. Piantadosi^{a,b}

^aUC Berkeley, Psychology ^bHelen Wills Neuroscience Institute

On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?

Emily M. Bender*
ebender@uw.edu
University of Washington
Seattle, WA, USA

Angelina McMillan-Major
aymm@uw.edu
University of Washington
Seattle, WA, USA

Timnit Gebru*
timnit@blackinai.org
Black in AI
Palo Alto, CA, USA

Shmargaret Shmitchell
shmargaret.shmitchell@gmail.com
The Aether

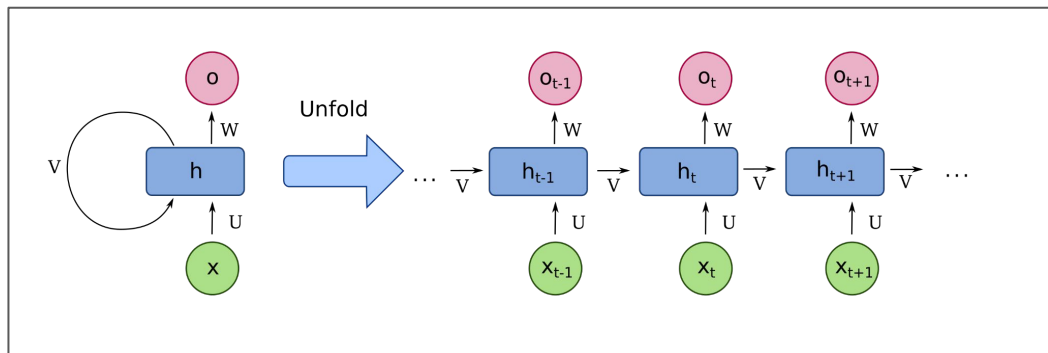
Lol, lmao even 

Hot (?) take: This is a contingent fact, not a necessary one

Why think about “Language”/Linguistics (in NLP)?

Linguistics can give useful insight!

- “Every time I fire a linguist, the performance of the speech recognizer goes up” (Frederick Jelinek, NLP pioneer)
- Counterpoint: RNNs! (Jeffrey Elman, *Finding Linguistic Structure in Time*)



Mamba: Linear-Time Sequence Modeling with Selective State Spaces

Albert Gu¹ and Tri Dao²

¹Machine Learning Department, Carnegie Mellon University

²Department of Computer Science, Princeton University
agu@cs.cmu.edu, tri@tridao.me

Google DeepMind

2024-3-1

Griffin: Mixing Gated Linear Recurrences with Local Attention for Efficient Language Models

Soham De¹, Samuel L. Smith¹, Anushan Fernando¹, Aleksandar Botev¹, George Cristian-Muraru¹, Albert Gu², Ruba Haroun¹, Leonard Berrada¹, Yutian Chen¹, Srivatsan Srinivasan¹, Guillaume Desjardins¹, Arnaud Doucet¹, David Budden¹, Yee Whye Teh¹, Razvan Pascanu¹, Nando De Freitas¹ and Caglar Gulcehre¹
¹Equal contributions, ¹Google DeepMind, ²Work done while at Google DeepMind

Language raises interesting problems

OpenAI Menu

Customer stories

Government of Iceland

How Iceland is using GPT-4 to preserve its language.



Dialect prejudice predicts AI decisions about people's character, employability, and criminality

Valentin Hofmann^{1-3*†}, Pratyusha Ria Kalluri⁴, Dan Jurafsky⁴, Sharese King^{5*}

¹Allen Institute for AI ²University of Oxford ³LMU Munich
⁴Stanford University ⁵The University of Chicago

How do you train a (trillion?) parameter LM to have good performance on low-resource languages?

How does linguistic variation affect model responses?
How *should* it?

LLMs Are jUsT sToChAsTiC pArRoTs

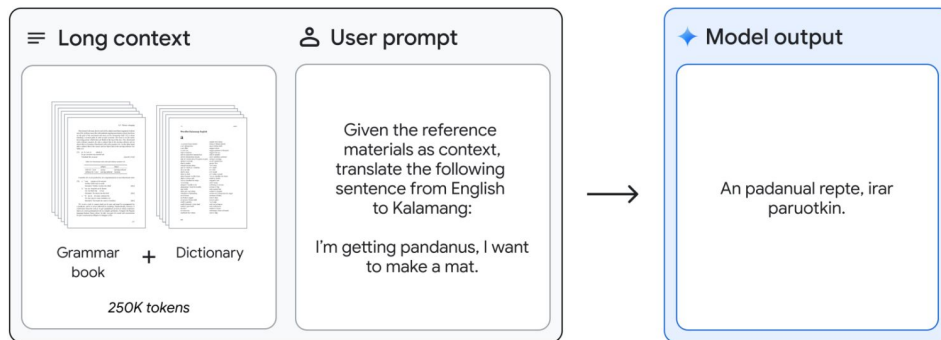


Figure 3 | Given a reference grammar book and a bilingual wordlist (dictionary), Gemini 1.5 Pro is able to translate from English to Kalamang with similar quality to a human who learned from the same materials.

Kalamang (200 native speakers, ~no digital training data)

Questions I'm interested in

Recall Poverty of the Stimulus argument from lecture:

Humans make consistent generalizations from very minimal, ambiguous data

=> Humans have innate bias towards “language” (whatever that looks like)

Two interesting points:

- Humans systematically choose A over B
- Humans learn to choose A over B with *very little data*

Q1: How do we get LLMs to make similar generalizations?

Suppose you want to turn natural language descriptions into code (method calls, etc).

How can you ensure that an LLM generalizes to:

- Arbitrarily deep function composition?
- Arbitrary (correct) variable identity?

These properties encapsulate *compositionality*

- This is a DEFINING property of natural language in general

Compositionality

jump

JUMP

turn around left

LTURN LTURN LTURN LTURN

jump thrice and turn left twice

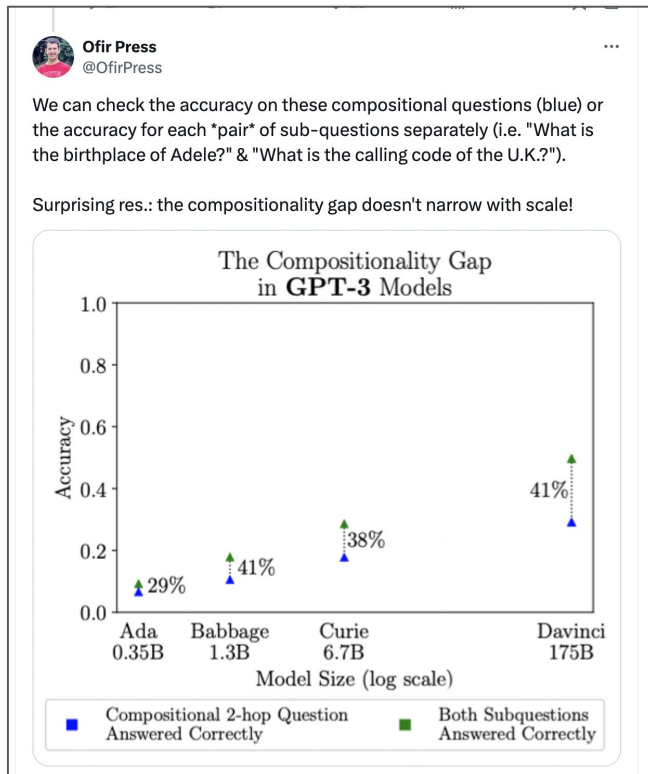
JUMP JUMP JUMP LTURN LTURN

jump opposite left after walk twice

WALK WALK LTURN LTURN JUMP

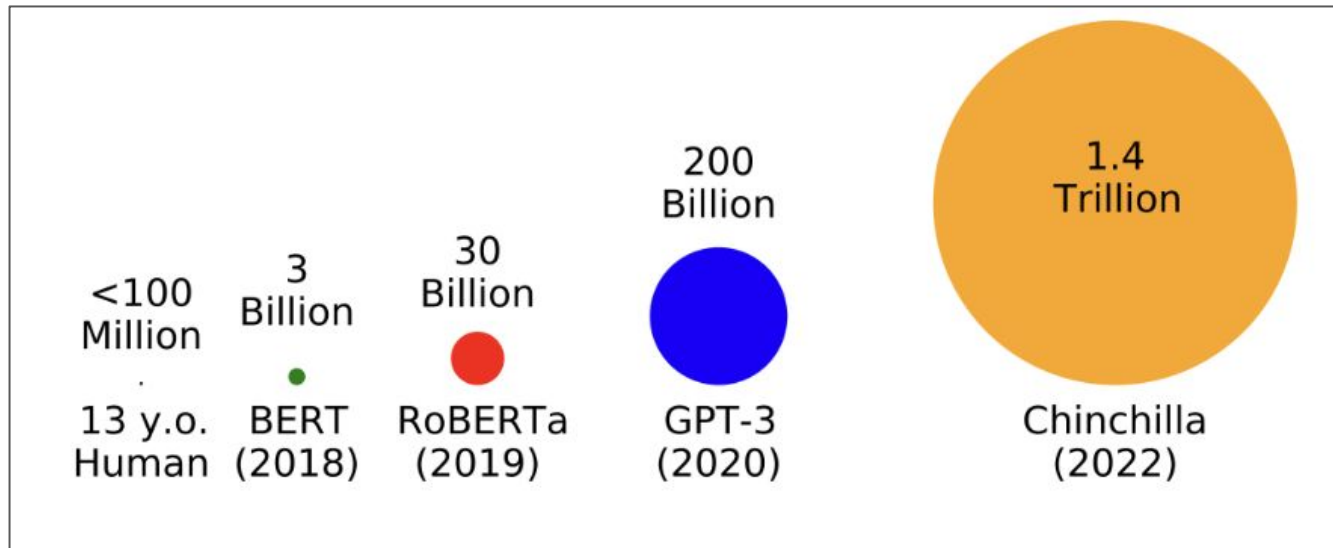
Figure 1: SCAN maps commands to actions

Small scale



Large scale

Q2: How do we get LLMs to be as efficient as people?



BabyLM challenge: train a language model on human-realistic data