Compositionality Without Classical Constituency

Raphaël Millière
Columbia University
rm3799@columbia.edu

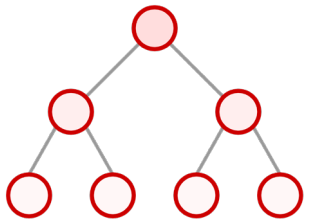

The Challenge of Compositionality for AI, 29-30 June 2022

The Principle of Compositionality

"The meaning of a whole is a function of the meanings of the parts and of the way they are syntactically combined" [Partee 1995]

- Underspecified (Whole/parts? Unique meaning? Strict determination?)
- Inadequate for natural language given background assumptions

The meaning of at least some complex expressions is at least partly determined by the meanings of their constituent parts + their syntactic structure

- Loosely extensible to other domains (e.g. images) *mutatis mutandis*

Compositional behavior

Given weak CP, we want systems that can process inputs compositionally

INPUT	OUTPUT
Man bites dog. Q: Who needs urgent care? A:	The dog

Compositional behavior

Given weak CP, we want systems that can process inputs compositionally

INPUT

A red cube on top of a blue cube

OUTPUT

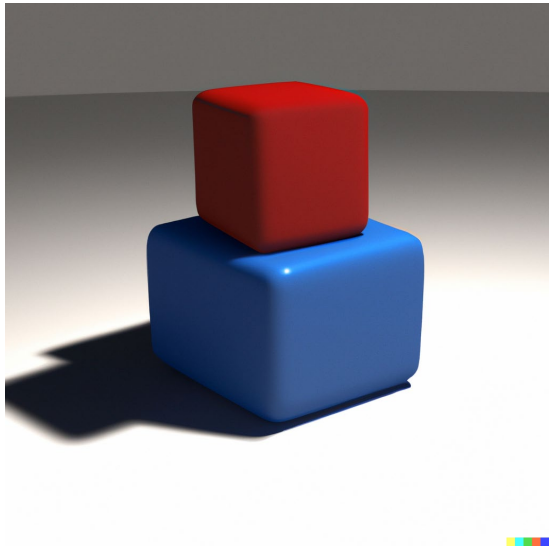

Compositional behavior

Given weak CP, we want systems that can process inputs compositionally

INPUT

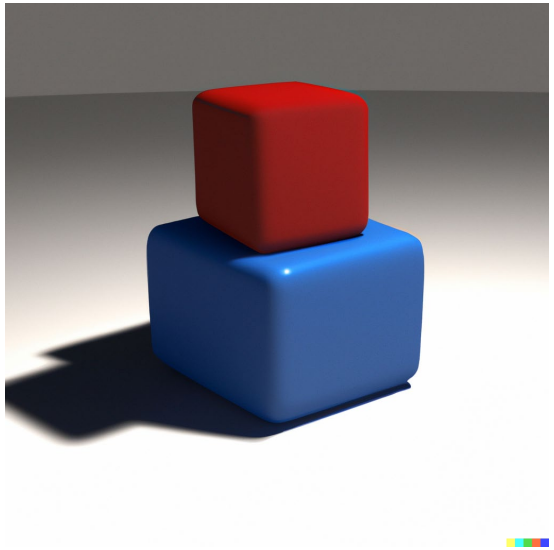

OUTPUT

A red cube on top
of a blue cube

Compositional behavior

Given weak CP, we want systems that can process inputs compositionally

INPUT

OUTPUT

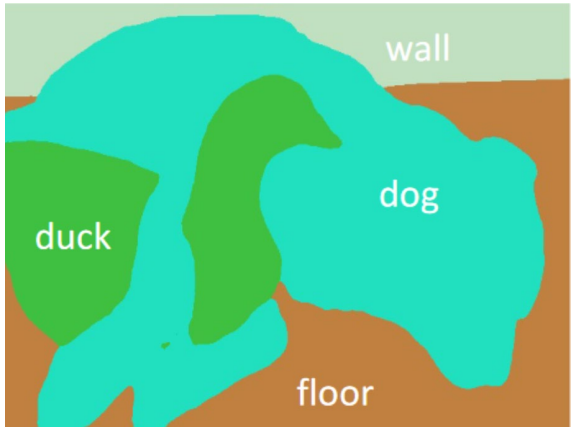

Yin et al. (2022)

Compositional representations

"Compositionality is the classic idea that new representations can be constructed through the combination of primitive elements" [Lake et al. 2016]

- Pertains to how representations are structured and processed
- 'Combination' can be understood in different ways

Classical constituency

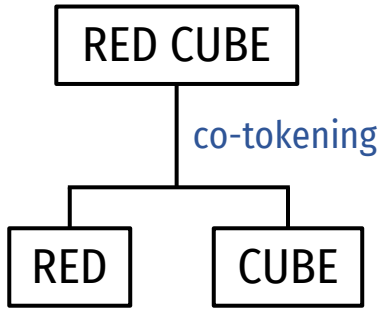

Derivation

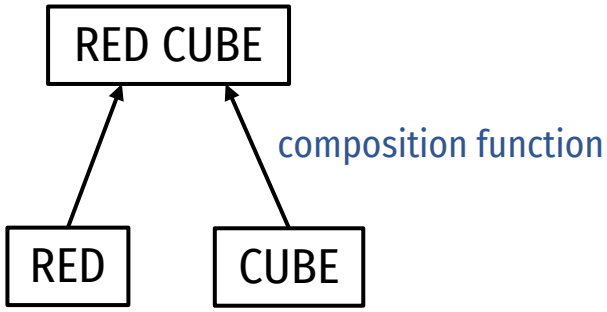

The theoretical challenge

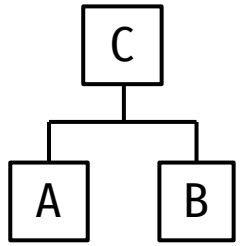

classical

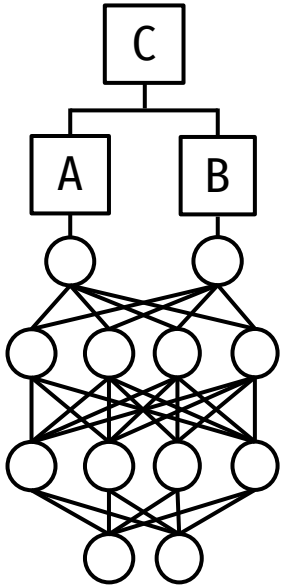

hybrid

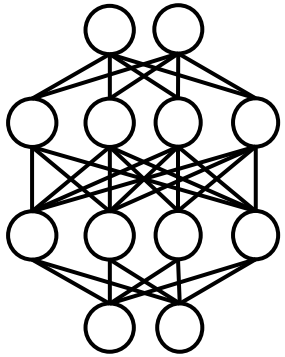

connectionist

- The old critique: connectionist systems lack a mechanism for combining representations compositionality [Fodor & Pylyshyn 1988]
- What have we learned 34 years on?

The methodological challenges

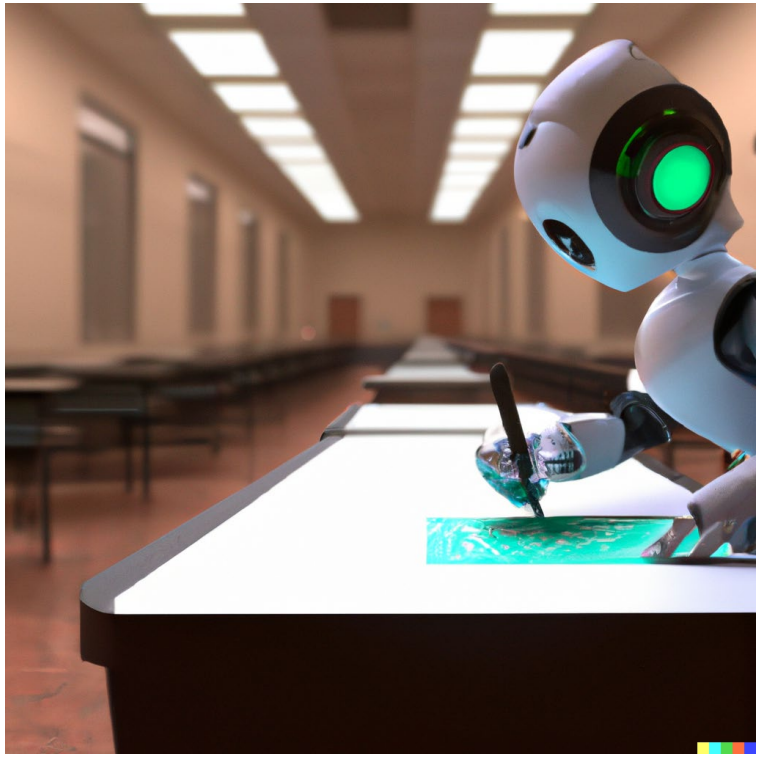

- How can we systematically evaluate compositional behavior?
- How can we get an insight at underlying mechanisms?

Conceptual combination

‘Syntax-light’ compositionality

- Requires compositing lexical meanings with minimal syntax + background knowledge [Ó Séaghdha 2008, Lake & Murphy 2021]

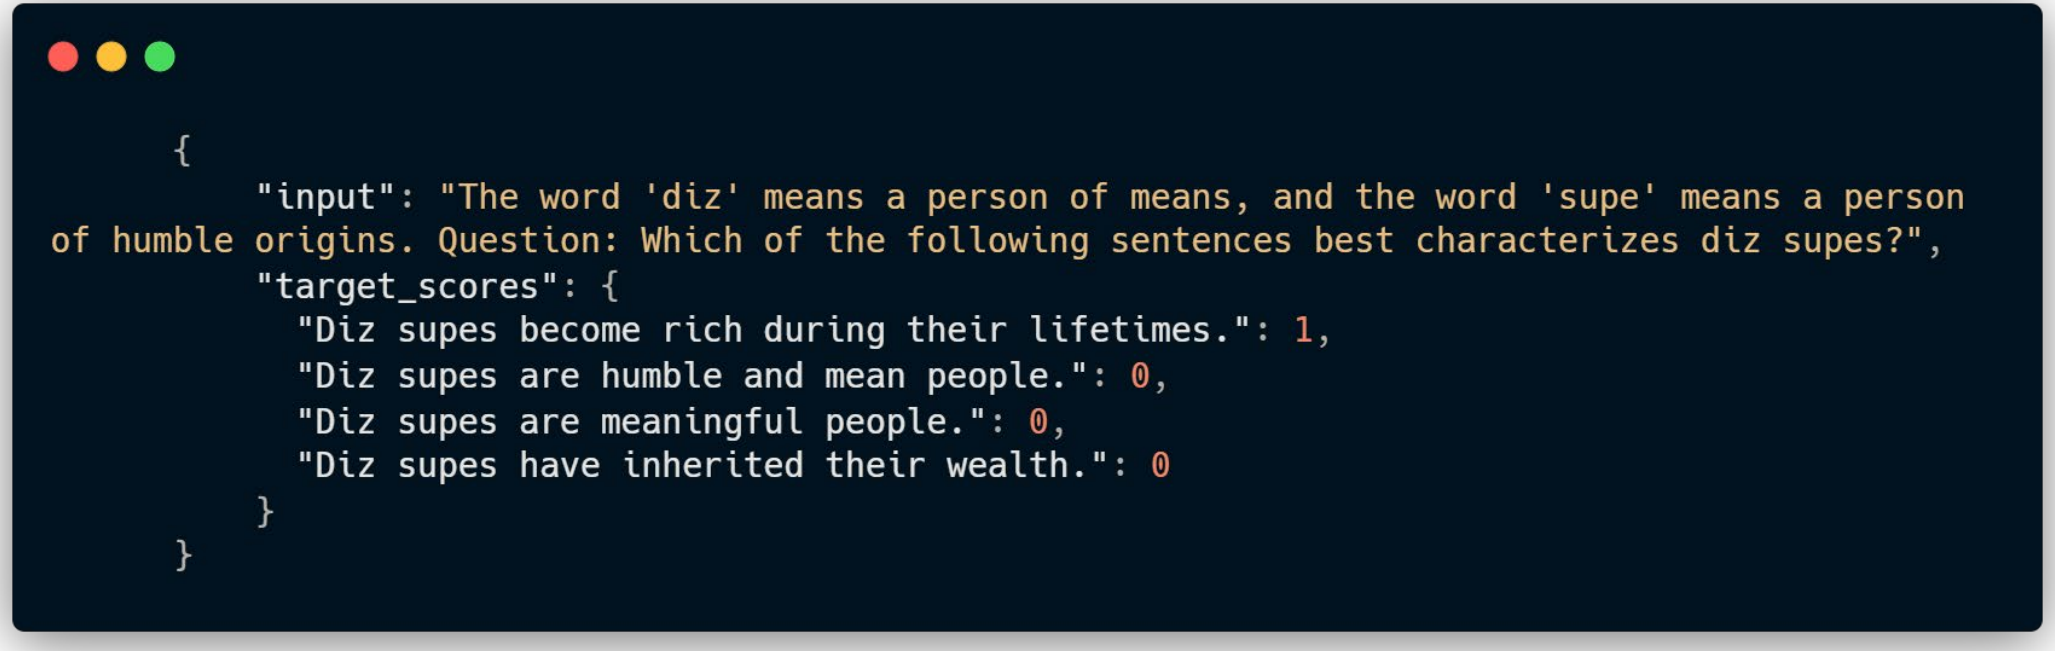

```
    {
        "input": "The word 'diz' means a person of means, and the word 'supe' means a person
of humble origins. Question: Which of the following sentences best characterizes diz supes?",
        "target_scores": {
          "Diz supes become rich during their lifetimes.": 1,
          "Diz supes are humble and mean people.": 0,
          "Diz supes are meaningful people.": 0,
          "Diz supes have inherited their wealth.": 0
        }
    }
```

BIG-bench conceptual_combinations task – Coelho Mollo, Millière, Rathkopf & Stinson in Srivastava et al. (2022)

Conceptual combination

‘Syntax-light’ compositionality

- Requires compositing lexical meanings with minimal syntax + background knowledge [Ó Séaghdha 2008, Lake & Murphy 2021]


```
    {
      "input": "Concept: burning questions. Question: Which of the following sentences best
characterizes burning questions?",
      "target_scores": {
        "Burning questions are urgent.": 1,
        "Burning questions are blazing hot.": 0,
        "Sunscreen can prevent burning questions.": 0,
        "Burning questions are multiple choice.": 0
      }
    }
```


BIG-bench conceptual_combinations task – Coelho Mollo, Millière, Rathkopf & Stinson in Srivastava et al. (2022)

Conceptual combination

‘Syntax-light’ compositionality

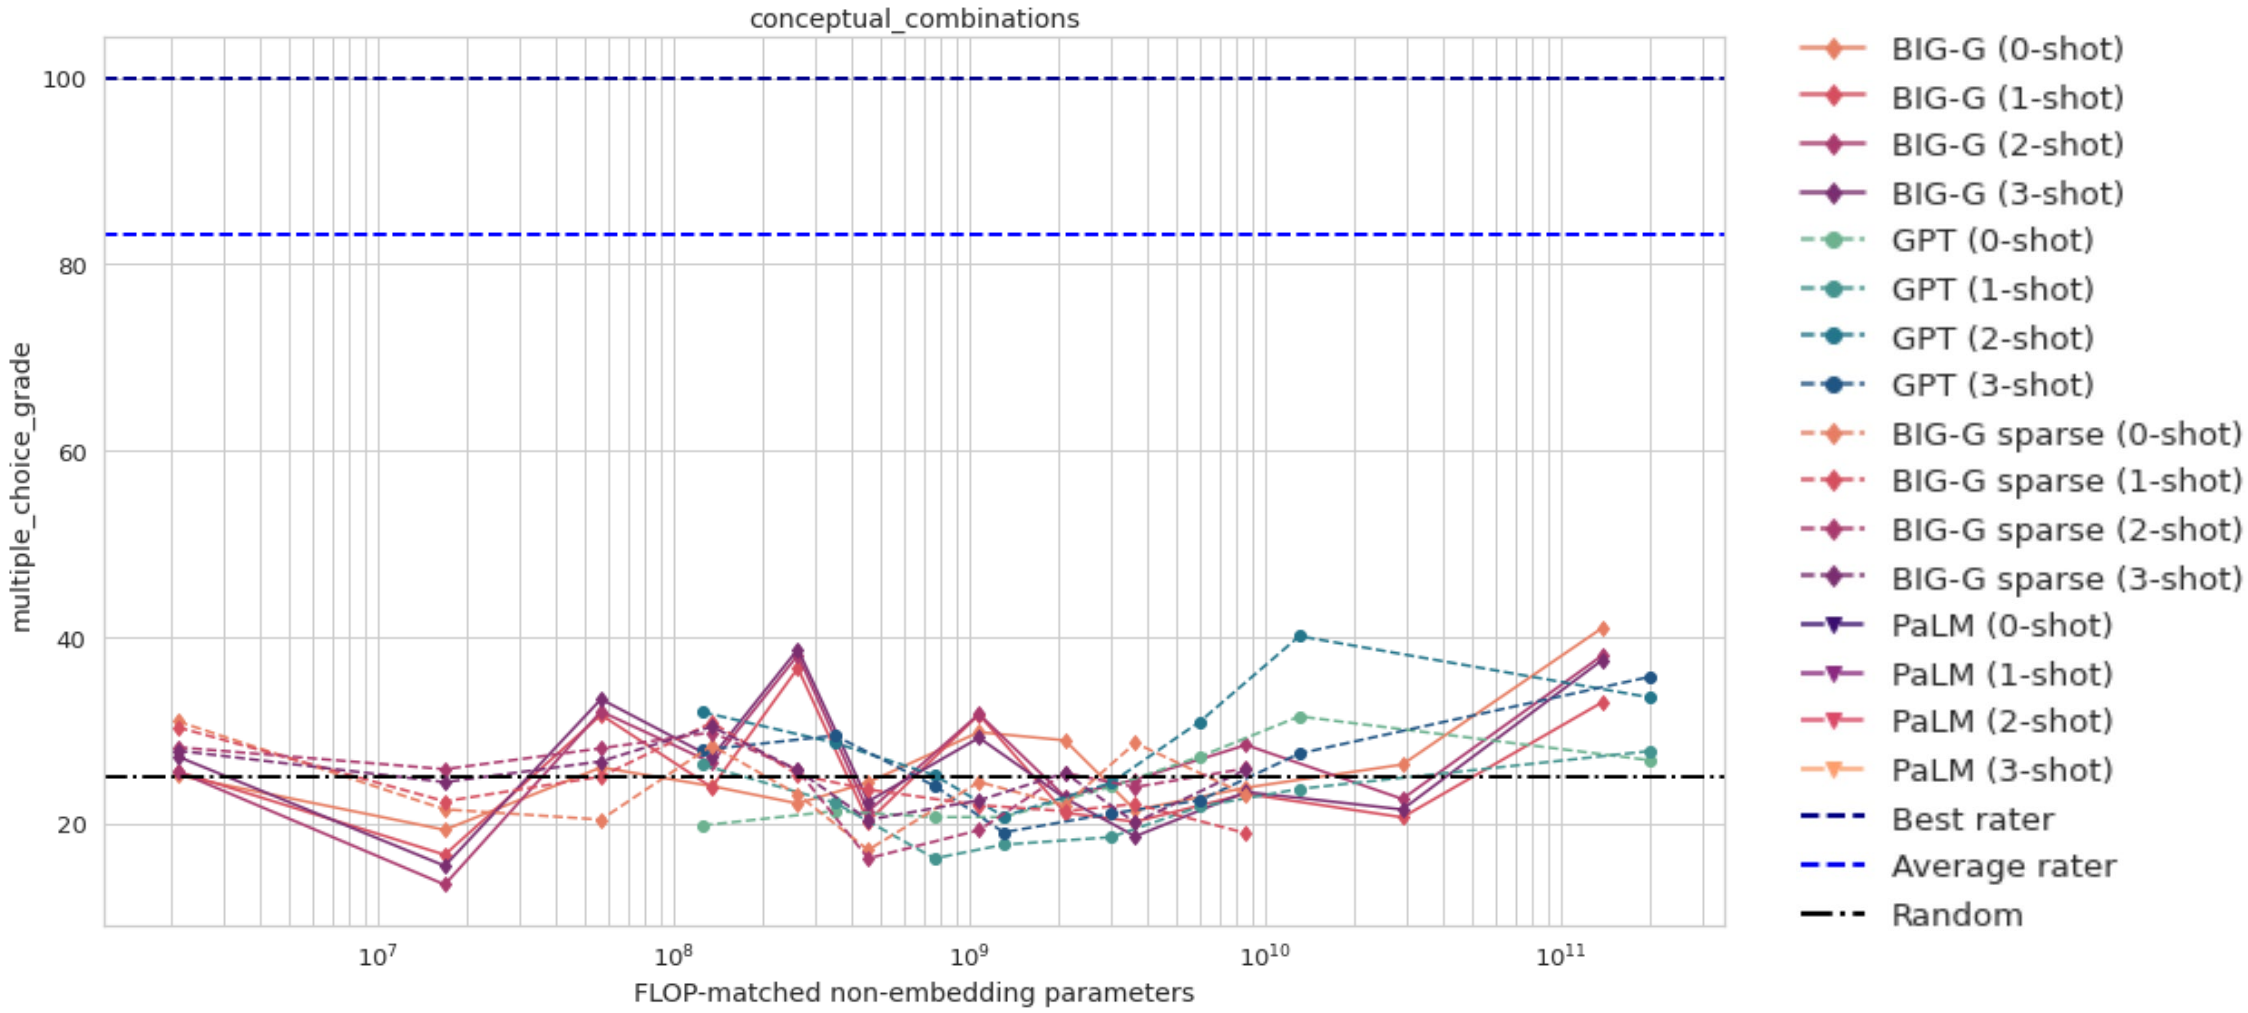

Conceptual combination

‘Syntax-light’ compositionality

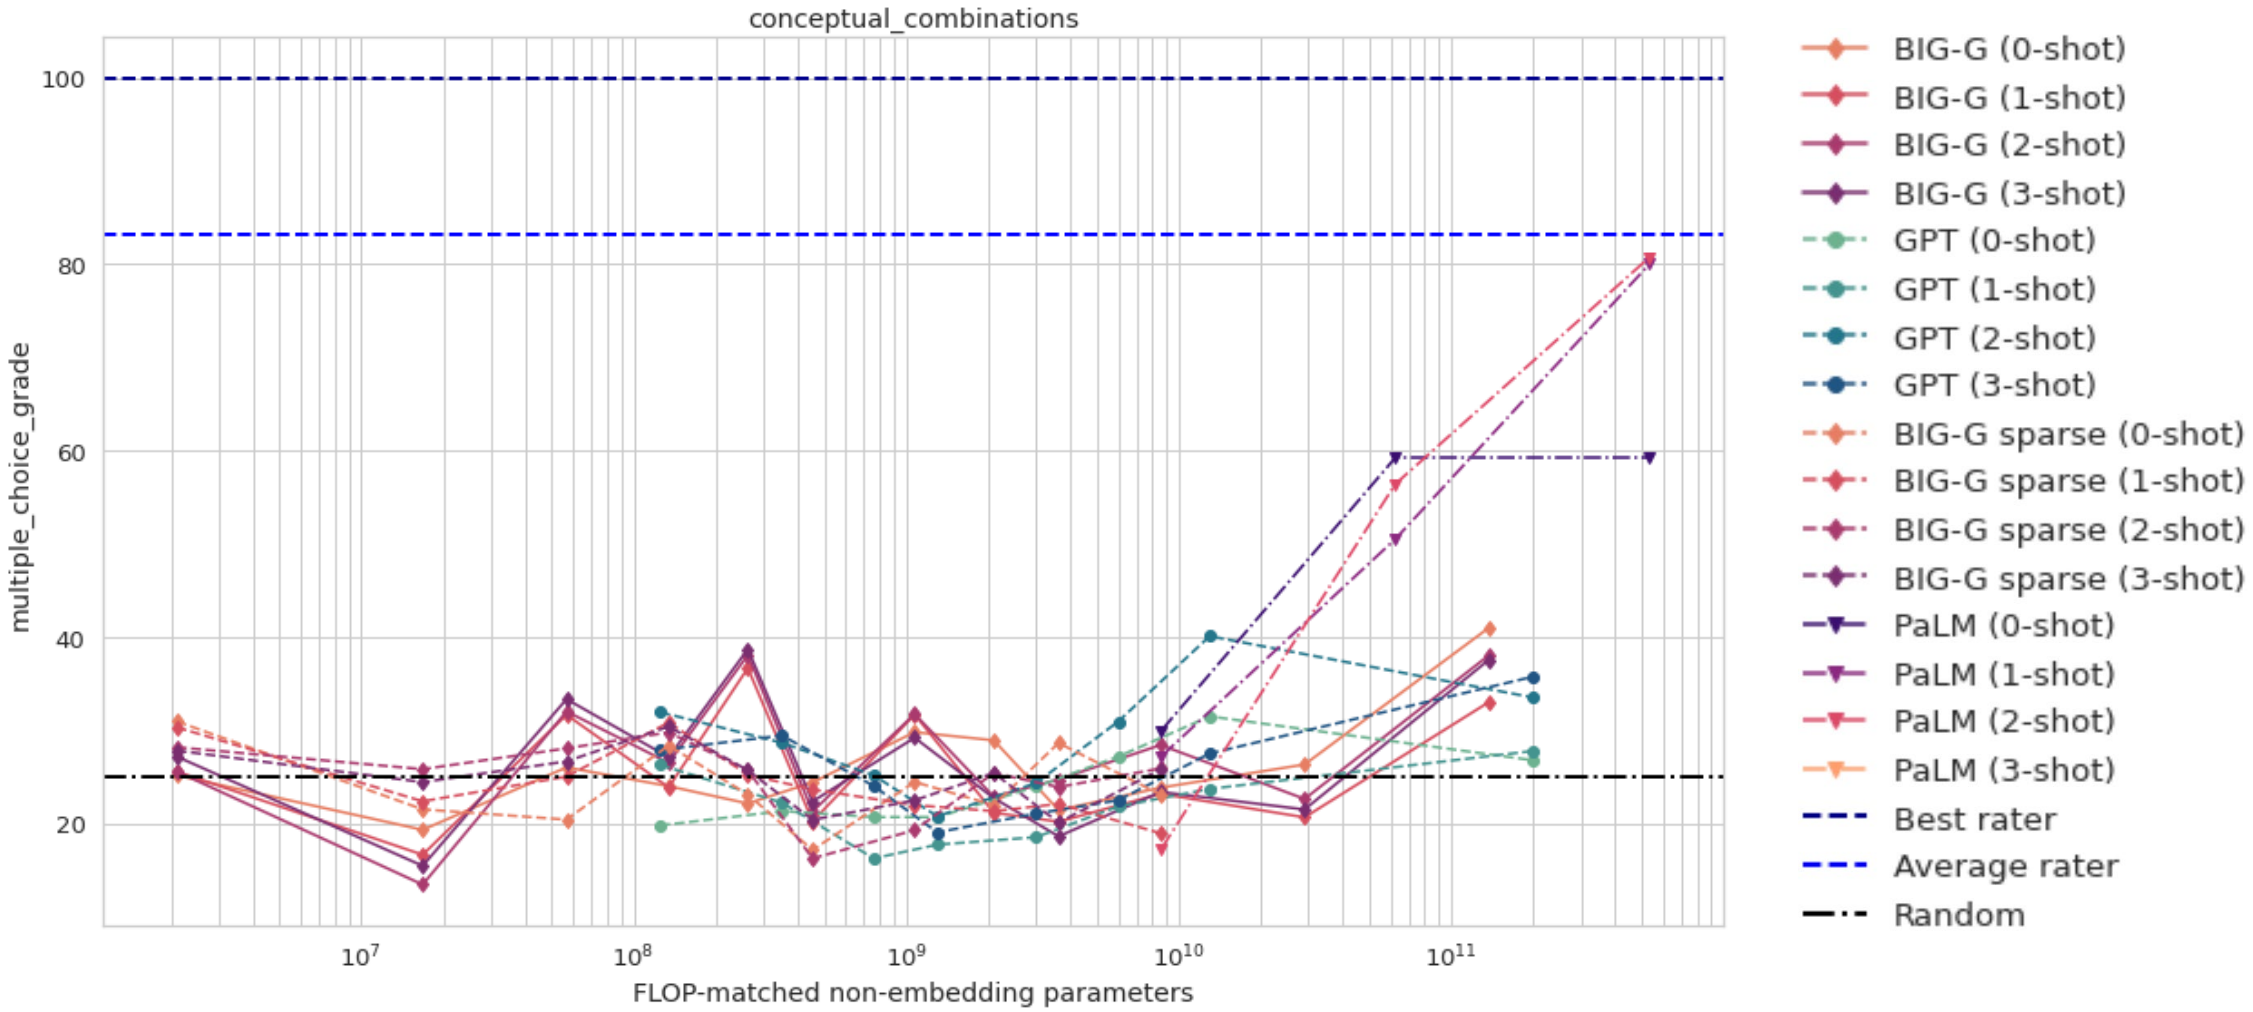

Conceptual combination

Across modalities

‘Shark Racoon’

‘Lion Squid’

‘Elephant Giraffe’

Made with DALL-E 2 [Ramesh et al. 2022]

Conceptual combination

Across modalities

‘The Great Wall of San Francisco’

‘The Great Wall of Stanford’

‘The Great Wall of Bali’

Made with Imagen [Saharia et al. 2022]
Source: David Ha

Complex compositional behavior

- Problem: no standard test for compositional behavior
- Different approaches:
  - SCAN [Lake & Baroni 2018]
  - PCFG [Hupkes et al. 2020]
  - COGS [Kim & Linzen 2020]
- Initial results were somewhat underwhelming (with caveats)
- But tweaks drastically improve performance [Ontañón et al. 2022, Csordás et al. 2022]
  - SCAN (length cutoff=26): 0.20 → **1.00**
  - PCFG productivity: 0.5 → **0.85**
  - PCFG systematicity: 0.72 → **0.96**
  - COGS: 0.35 → **0.80**

Complex compositional behavior

DALL-E 2 [Ramesh et al. 2022]

Parti [Yu et al. 2022]

"A wombat sits in a yellow beach chair, while sipping a martini that is on his laptop keyboard. The wombat is wearing a white panama hat and a floral Hawaiian shirt. Out-of-focus palm trees in the background."

The devil is in the details

- Architecture / inductive biases
  - Self-attention is a very powerful mechanism
  - Minor tweaks go a long way
- Learning objective
  - Masked modeling superior to contrastive learning (?)
- Model size
  - Scaling laws + emergent capabilities [Wei et al. 2022]
- Dataset
  - Data structure [Chan et al. 2022, Akyürek & Andreas 2022]

Towards mechanistic understanding

- Transformers can implement complex algorithms, using compositions of attention heads, that keep track of compositional structure
- They induce a repertoire of non-content-specific computations [Shea 2021]
  - Independent of the particular values of the representations taken as input
- They approximate or 'soft' variable binding by dynamically routing information in layer subspaces acting as memory
  - Cf. *1G neurocompositional computing* [Smolensky et al. 2022]
- This goes a long towards explaining compositional behavior and in-context learning
  - Further investigation is needed to track specific computations

Edelman et al. 2021, Elhage et al. 2021, Weiss et al. 2021, Olsson et al. 2022, Elhage et al. 2022

Let's play a game. I will show you some variable assignments. Then I'll show you a sequence of variables. And you have to output the corresponding sequence of letters

Sure, that sounds fun.

Here's an example
Assignments:
x = a
y = d
z = p
w = k
Sequence: y z x w
Answer: d p a k

Got it!

Assignments:
x = a
y = d
z = p
w = k
Sequence: w x y z
Answer?

k a d p

Assignments:
x = v
y = t
z = h
w = g
Sequence: x w z y
Answer?

v g h t

Let's try a variation

Sure!

Assignments:
Z = fred
V = mary
W = jackie
X = henry
Sequence: V W X Z
Answer?

mary jackie henry fred

Chinchilla [Hoffmann et al. 2022]
Source: Murray Shanahan

Conclusion

- DNNs can be given the resources to behave compositionally by inducing suitable composition functions *if* they have the right features (biases, objective, size, data)
- Compositionality does *not* require constituent structure
- Relinquishing constituency comes with significant challenges and benefits
- Lessons for cognitive science?

Constituency

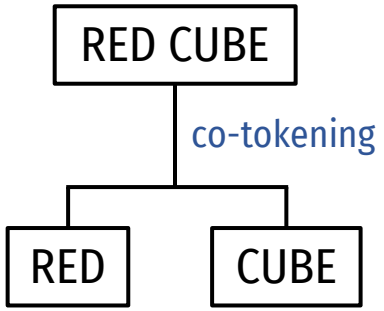

Derivation

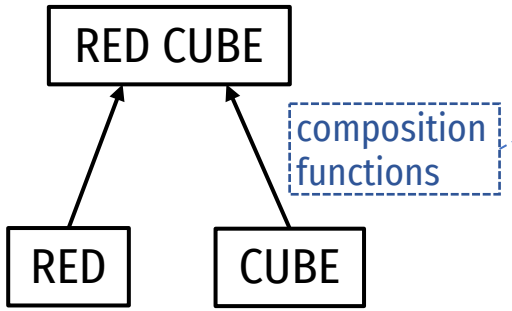

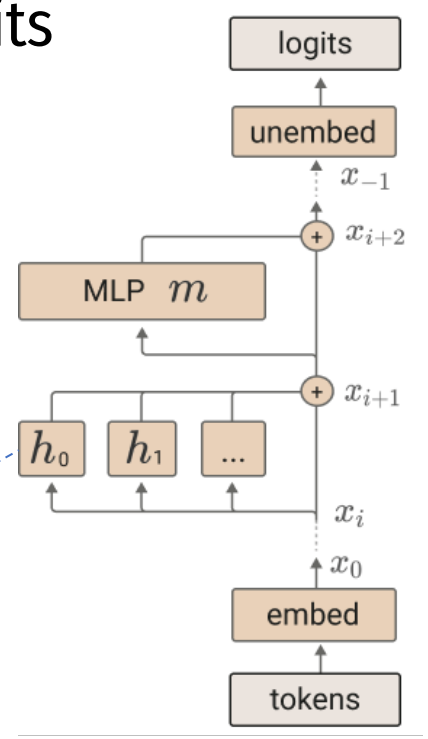

Thank you!

rm3799@columbia.edu

Visual compositionality

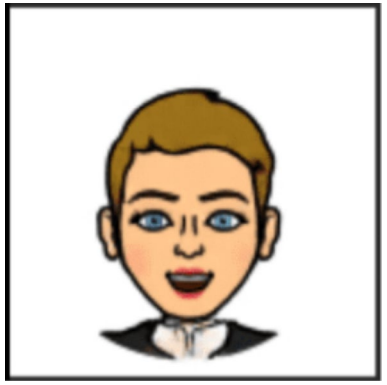

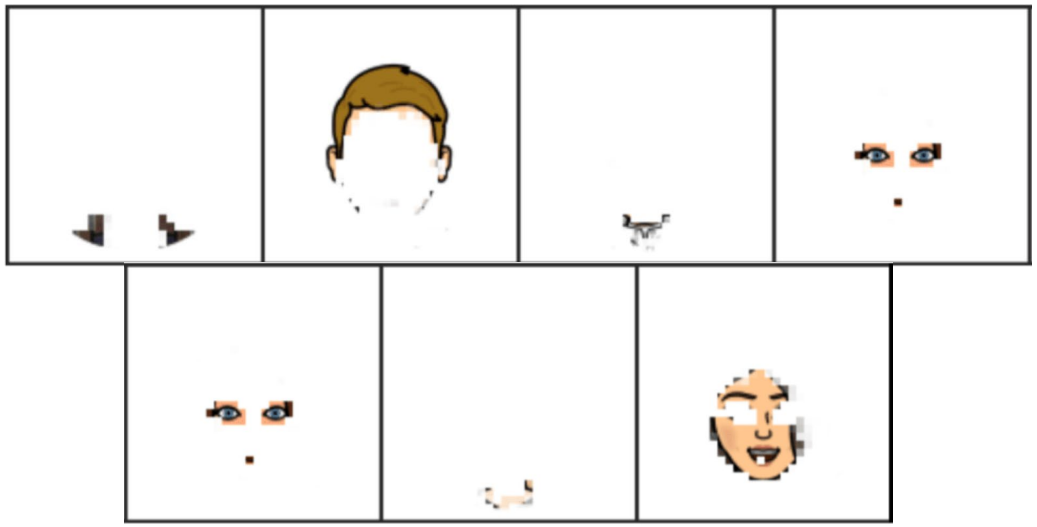

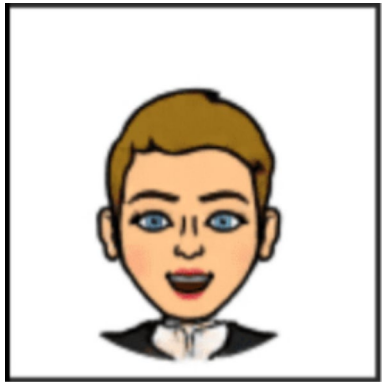

Singh et al. (2022)

“What makes [mental representations] compositional is that the content of structurally complex mental symbols is inherited from the contents of their less structurally complex parts” [Fodor 1997]

“There is no complete, precise formal account of the construction of composites or of mental processes in general that can be stated solely in terms of context-independent semantically interpretable constituents” [Smolensky 1995]