Compositionality in NLP

Generalisation and Interpretability with DisCoCirc

Tiffany Duneau

Quantinuum, University of Oxford

Andrew moved to the kitchen.
Clara journeyed to the park.
Andrew picked up the slippers.

⋮

Is Andrew in the park?

Grammar

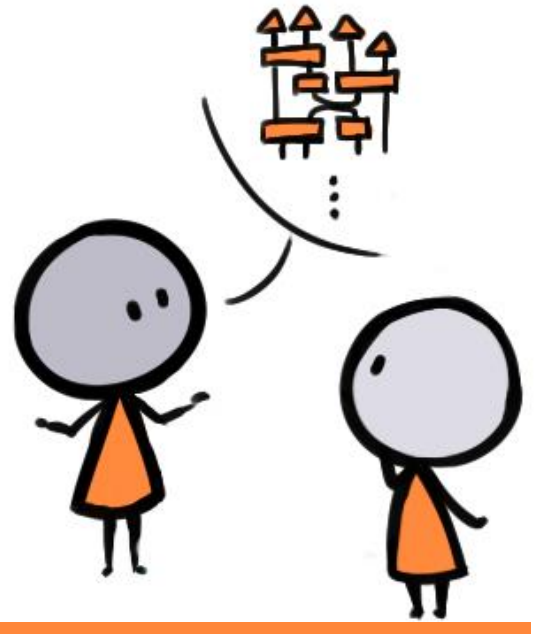

DisCoCirc

Wang-Mascianica et al, 2023. arXiv: 2301.10595

A monoidal category

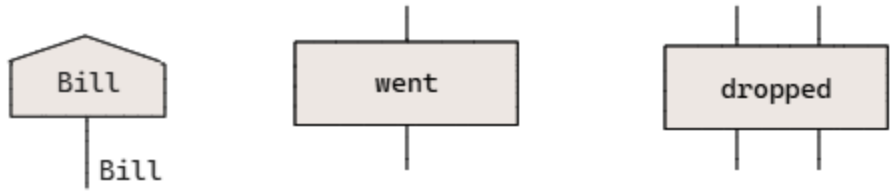

Noun

Intransitive Verb (ITV)

Transitive Verb (TV)

A monoidal category (ish)

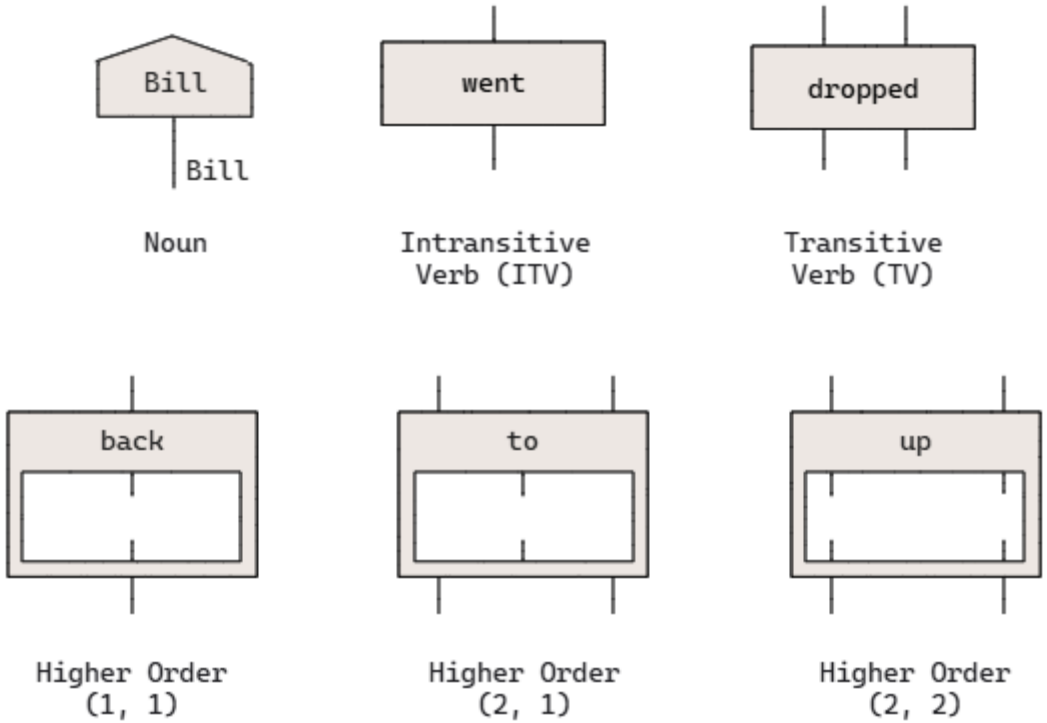

...where diagrams are sentences

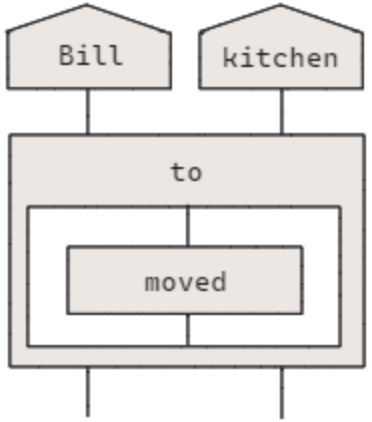

Bill moved to the kitchen.

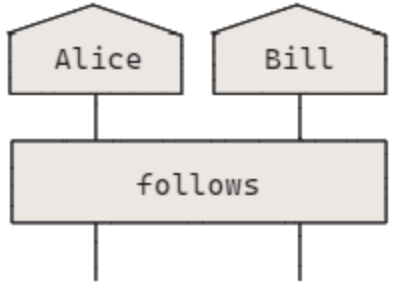

Alice follows Bill.

...that can compose into stories

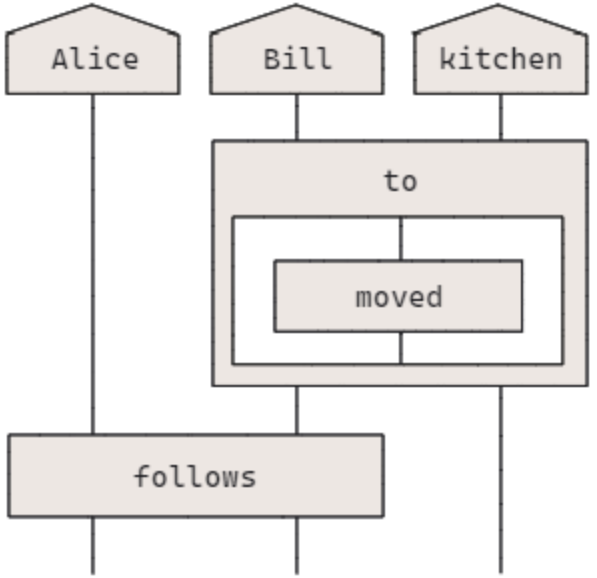

Bill moved to the kitchen.
Alice follows Bill.

Compositionality

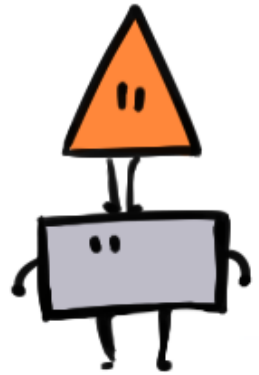

The meaning of the whole...

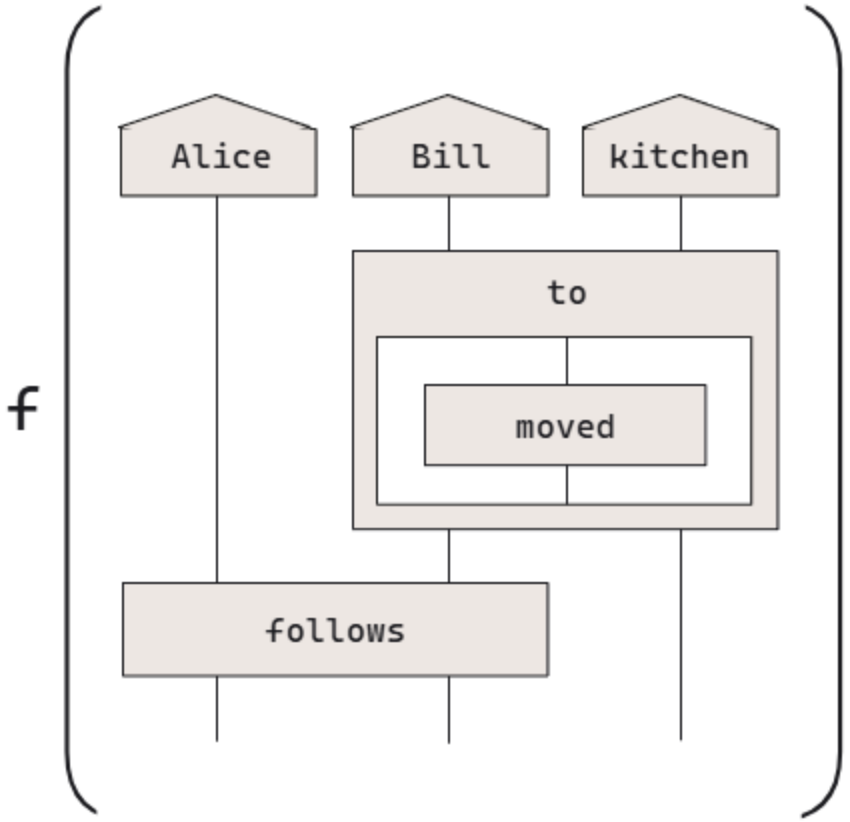

...depends on the meaning of the parts

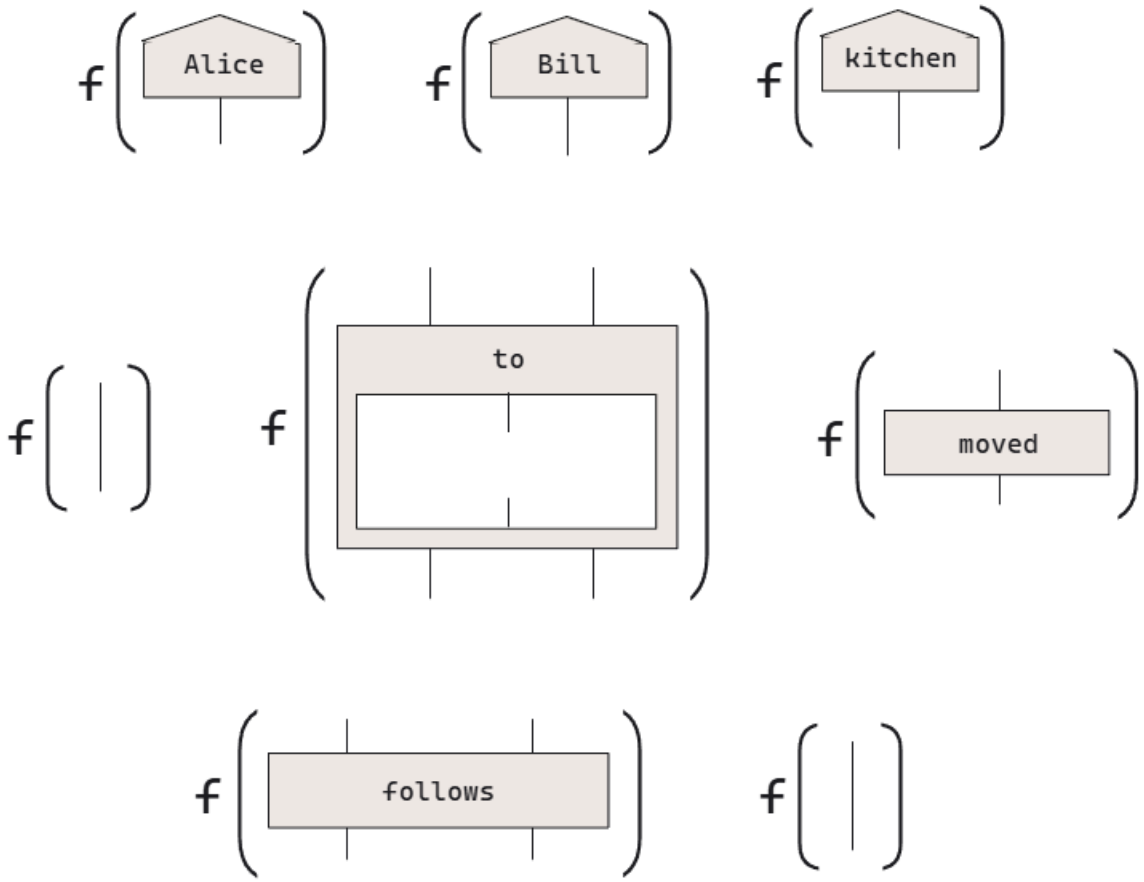

...and how they are put together

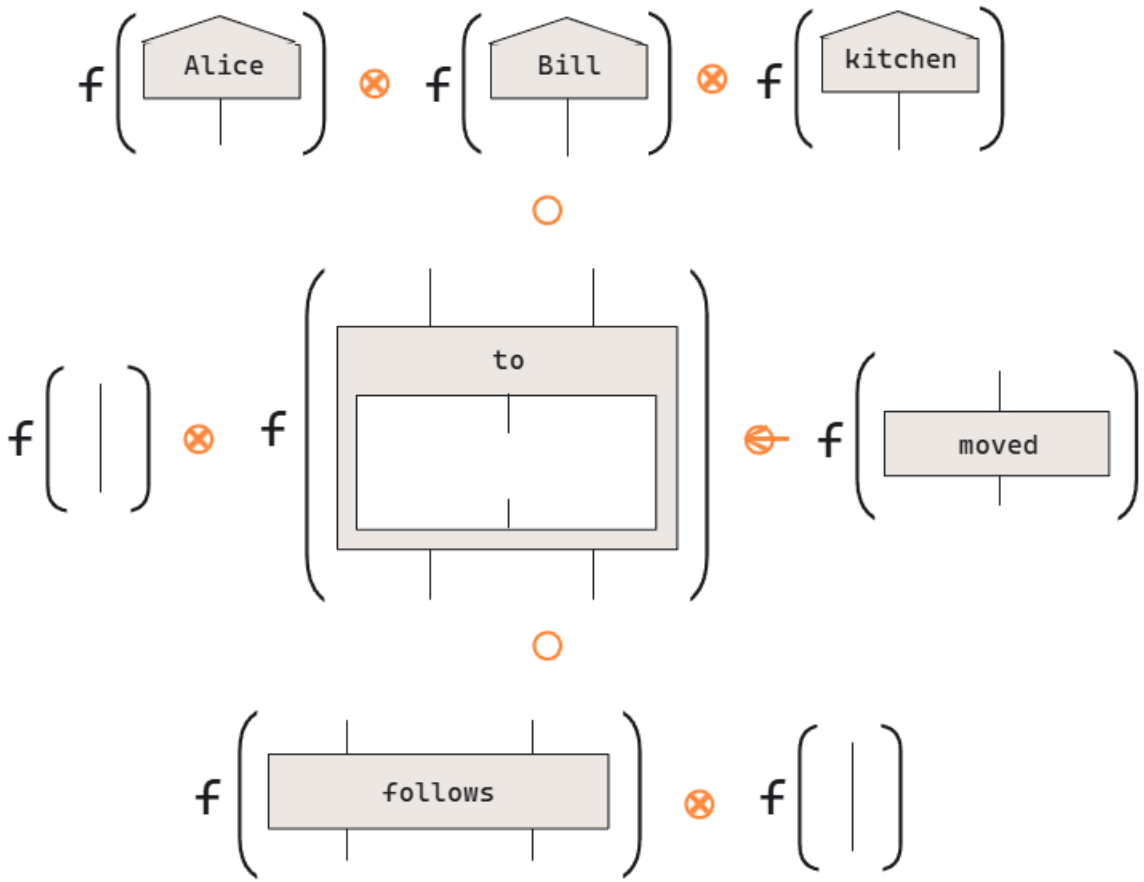

Testing Compositionality

Productivity

Hupkes et al, 2020. arXiv:1908.08351

Productivity

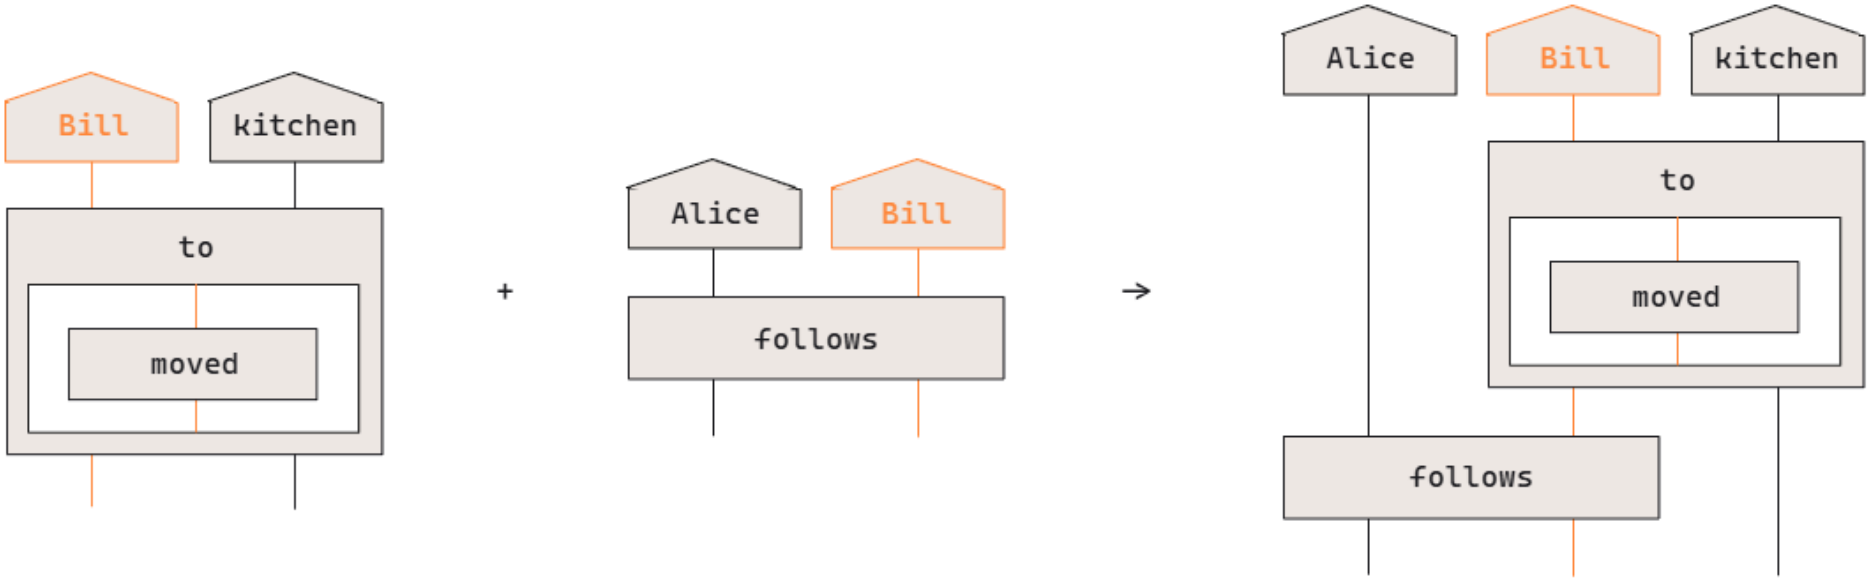

Compositionality Score

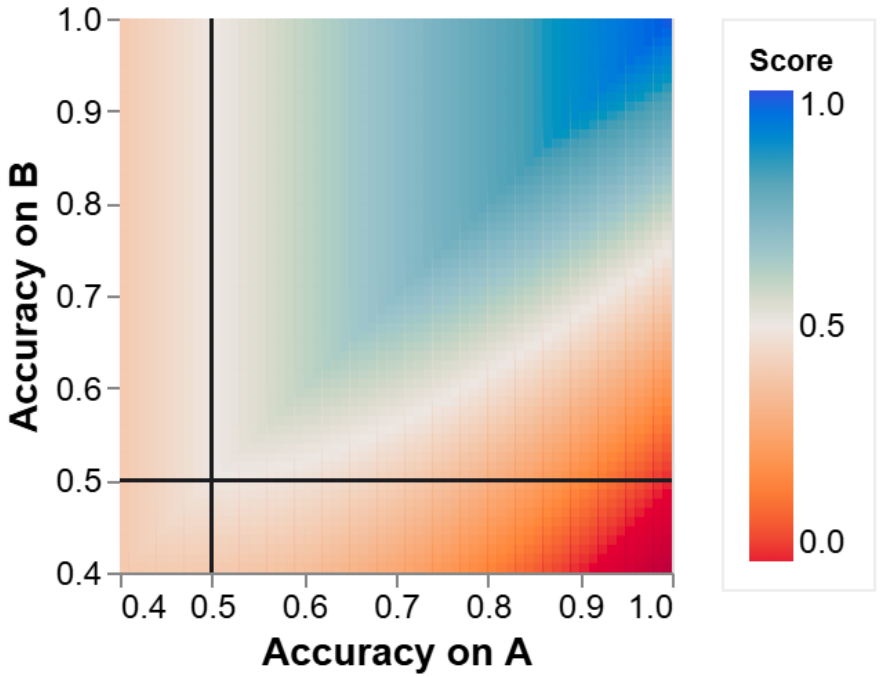

Data split

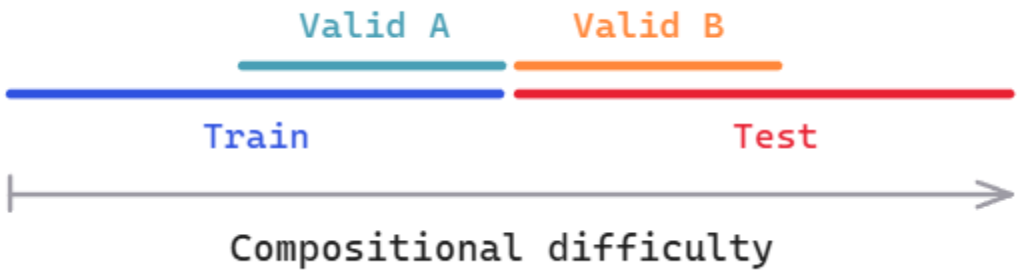

Task

bAbi 6

Weston et.al, 2015. arXiv:1502.05698

Question Answering

Bill moved to the kitchen.
Bill picked up the slippers.

Is Bill in the kitchen?

Assertions

Bill moved to the kitchen.
Bill picked up the slippers.

Is Bill in the kitchen?

↓

Bill is in the kitchen.
Bill is not in the kitchen

→

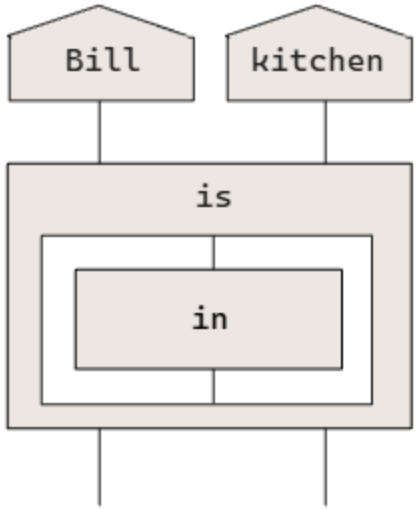

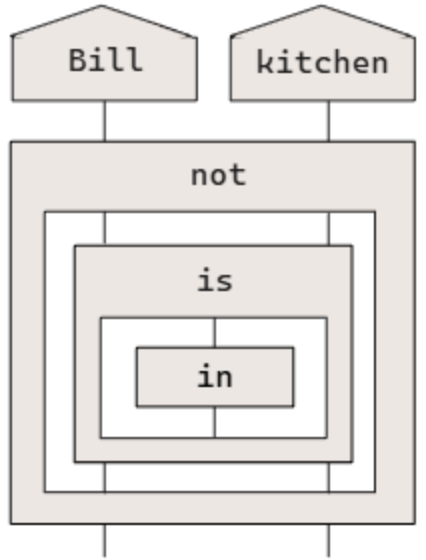

Assertions

Bill moved to the kitchen.
Bill picked up the slippers.

Is Bill in the kitchen? →

↓

Bill is in the kitchen.
Bill is not in the kitchen

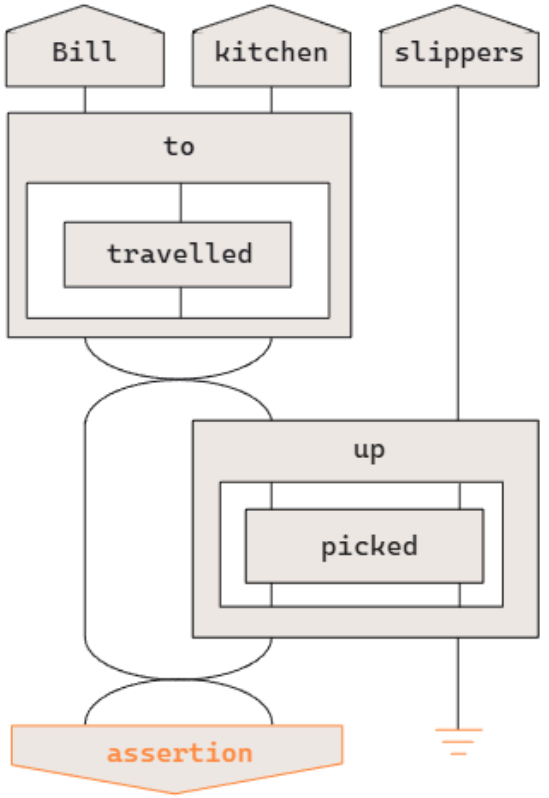

Models

Quantum

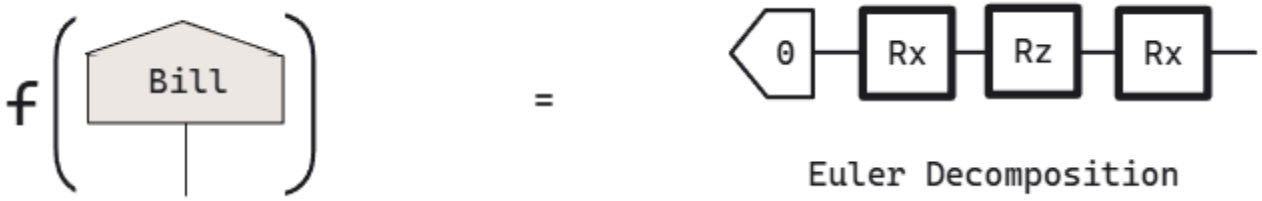

Quantum

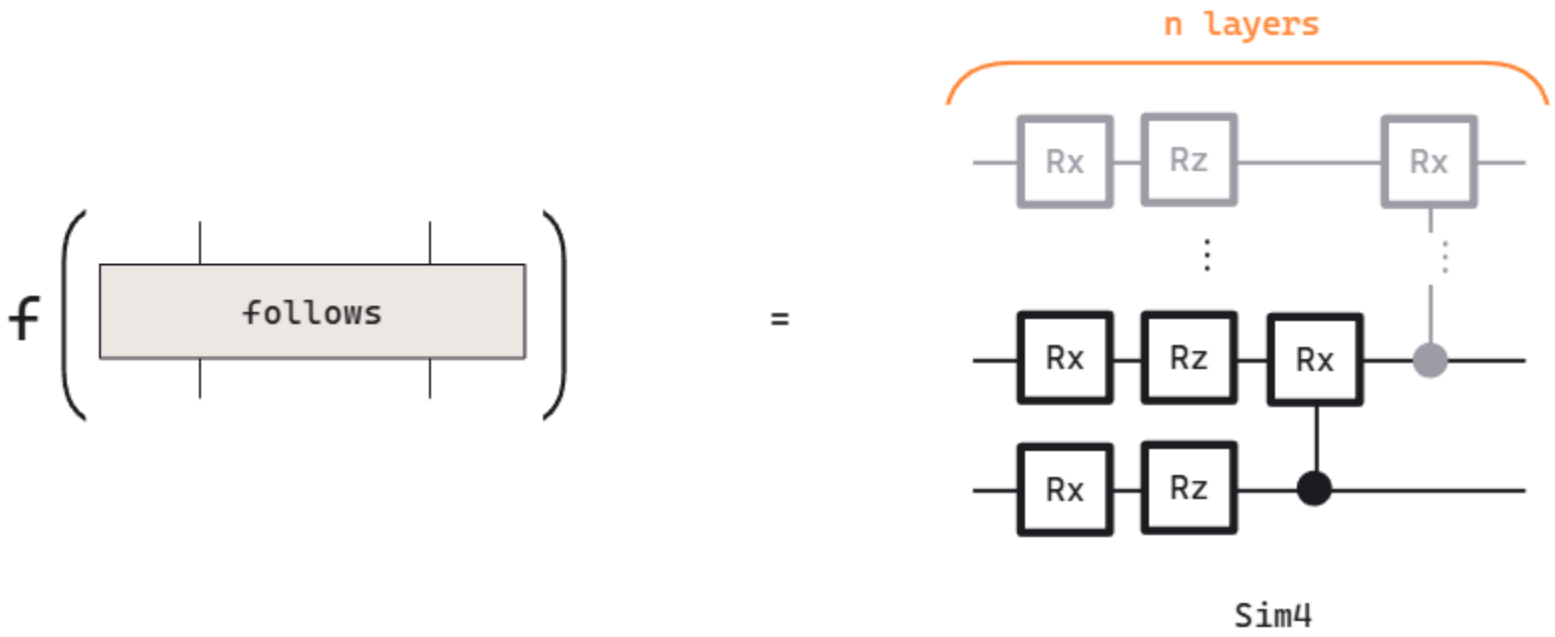

Sim et al, 2019. arXiv:1905.10876

Sandwich

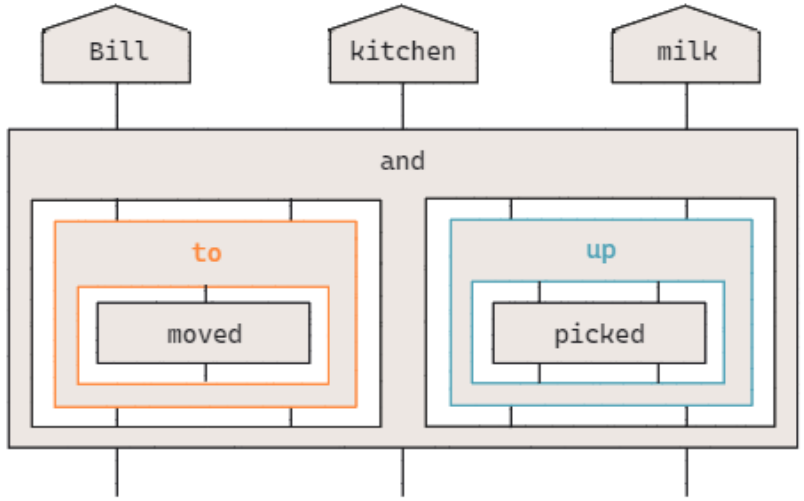
Bill
kitchen
milk
and
to
moved
up
picked

→

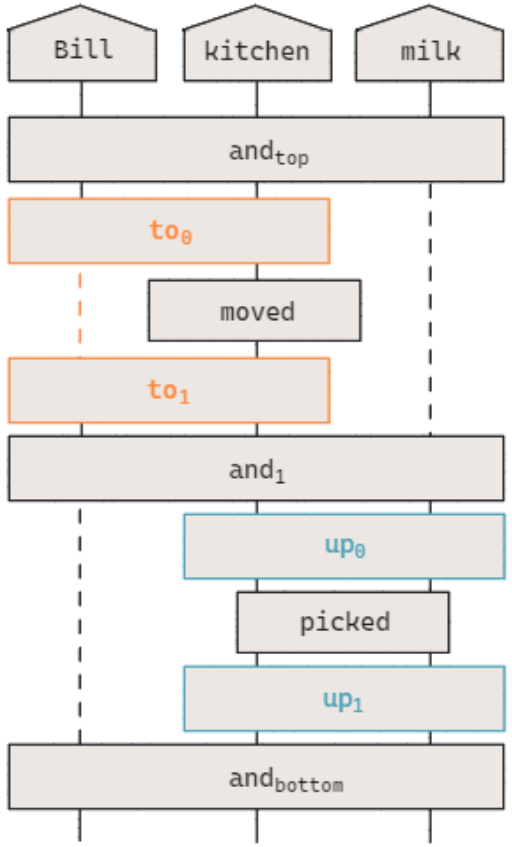
Bill
kitchen
milk
and_top
to_0
moved
to_1
and_1
up_0
picked
up_1
and_bottom

Neural

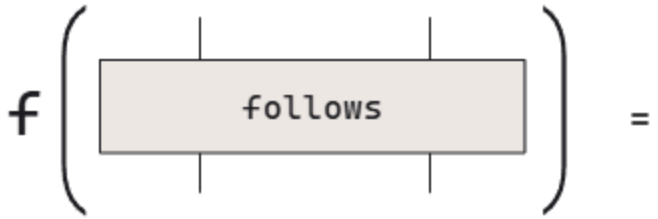

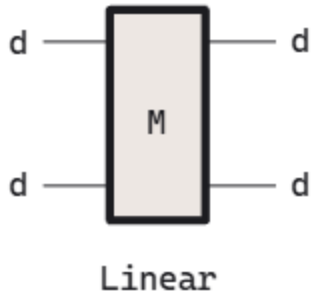

Neural

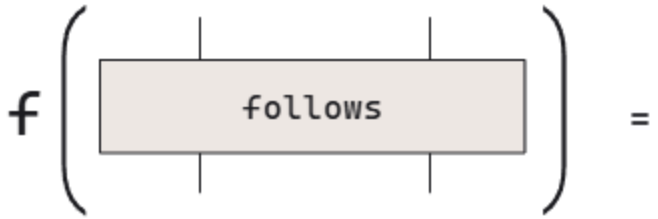

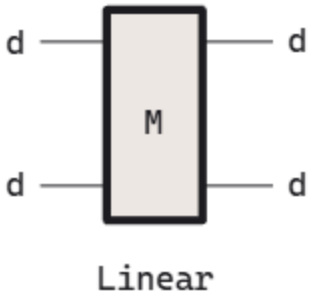

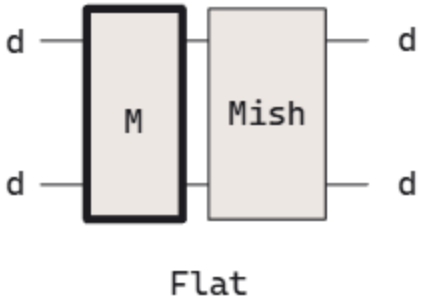

Neural

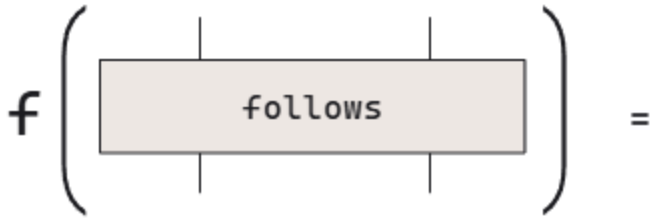

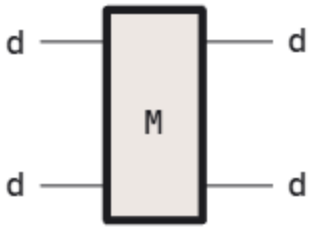

Linear

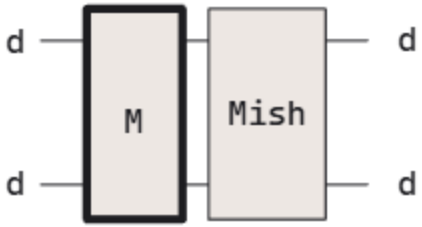

Flat

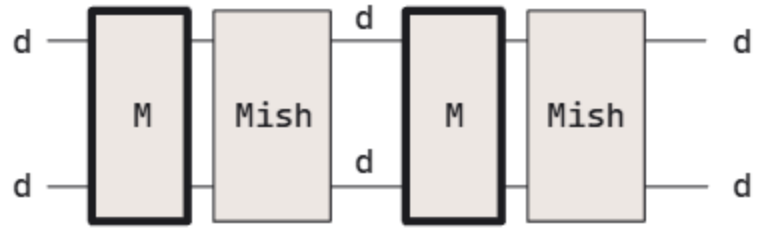

Hidden(1)

Compositional Generalisation

"Train small, test big"

Reading the plots

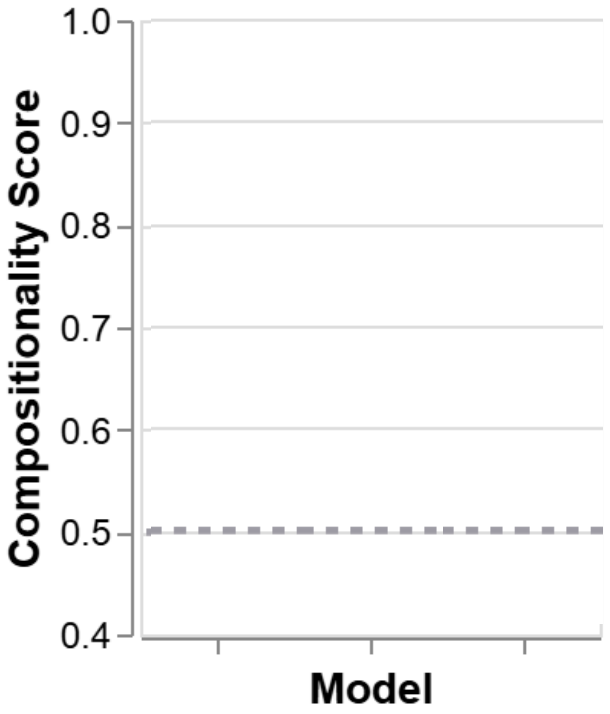

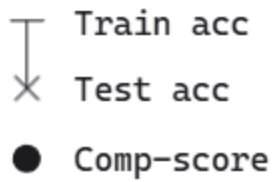

Proof-of-Concept: Following

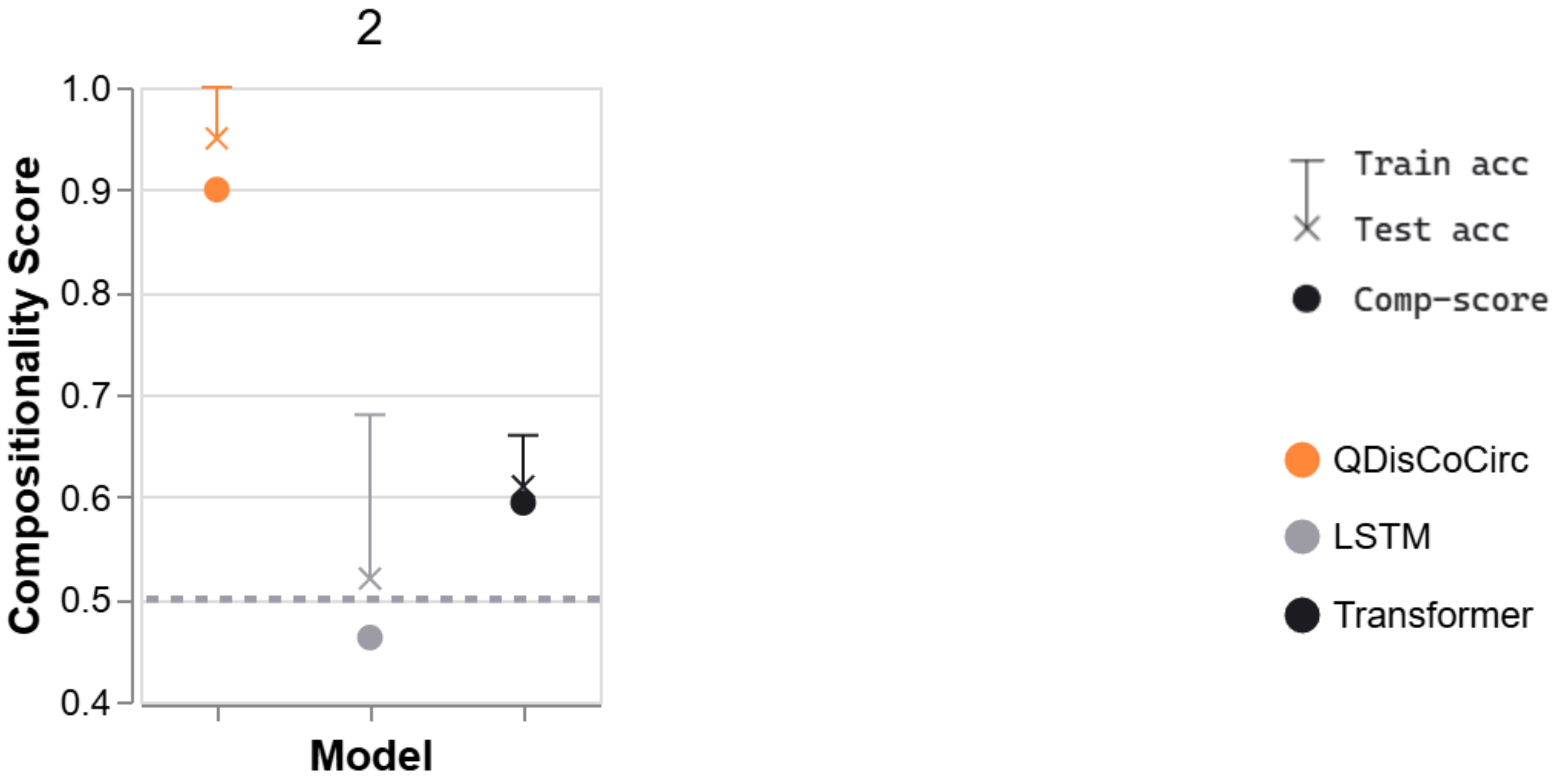

Duneau et al, 2024. arXiv:2409.08777

Proof-of-Concept: Following

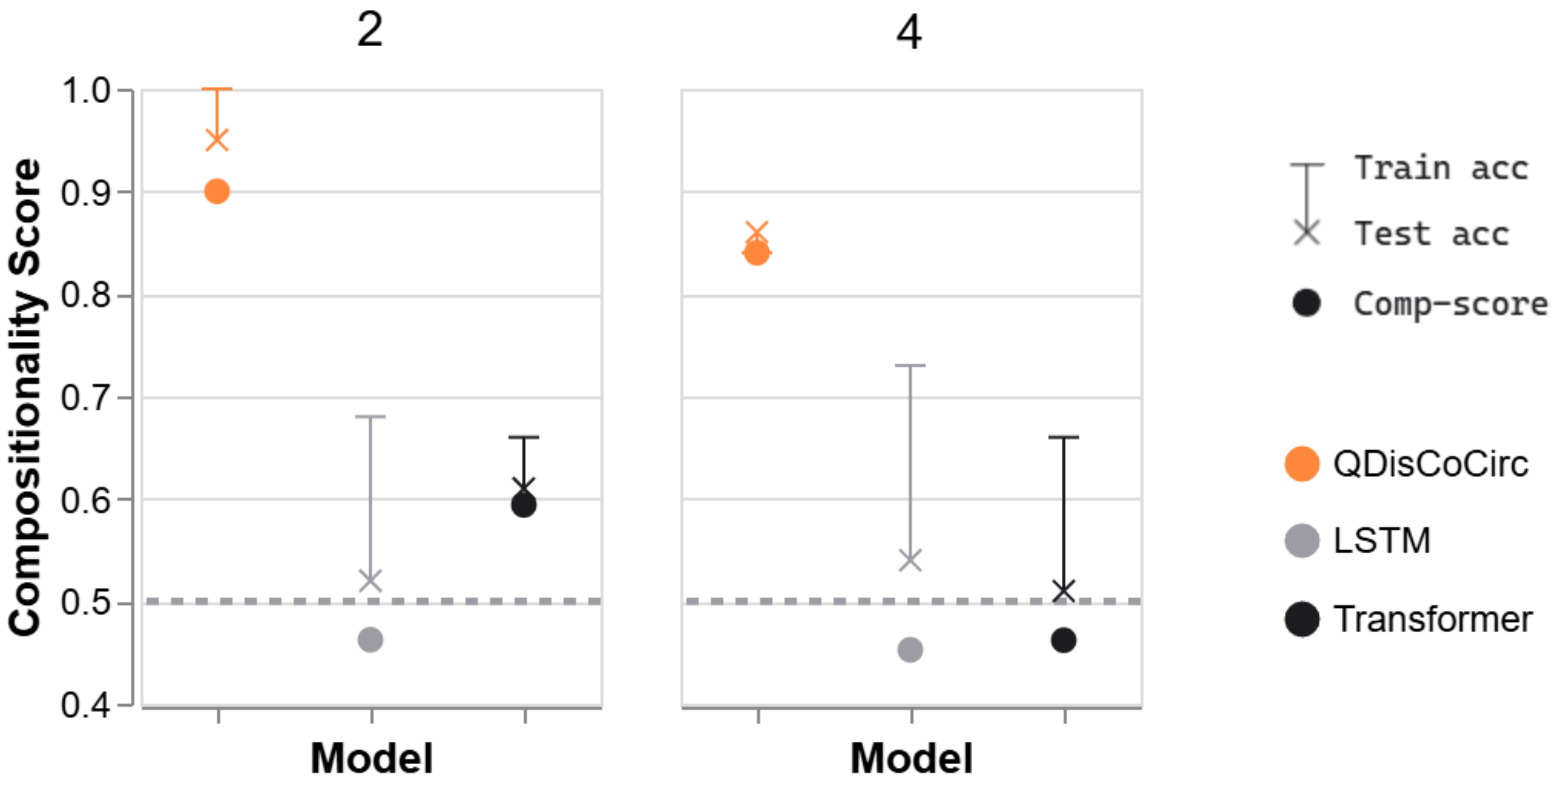

Duneau et al, 2024. arXiv:2409.08777

Productivity: bAbI 6

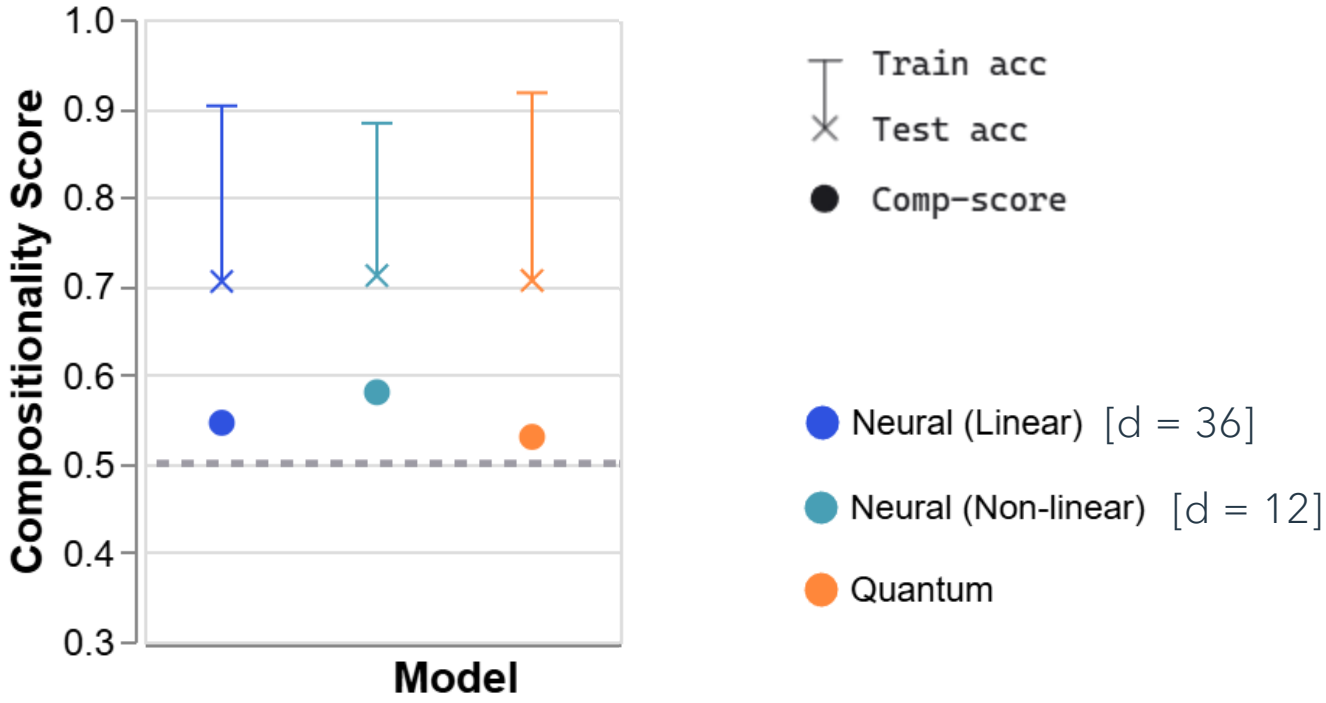

Compositional Interpretability

Black box

Many black boxes

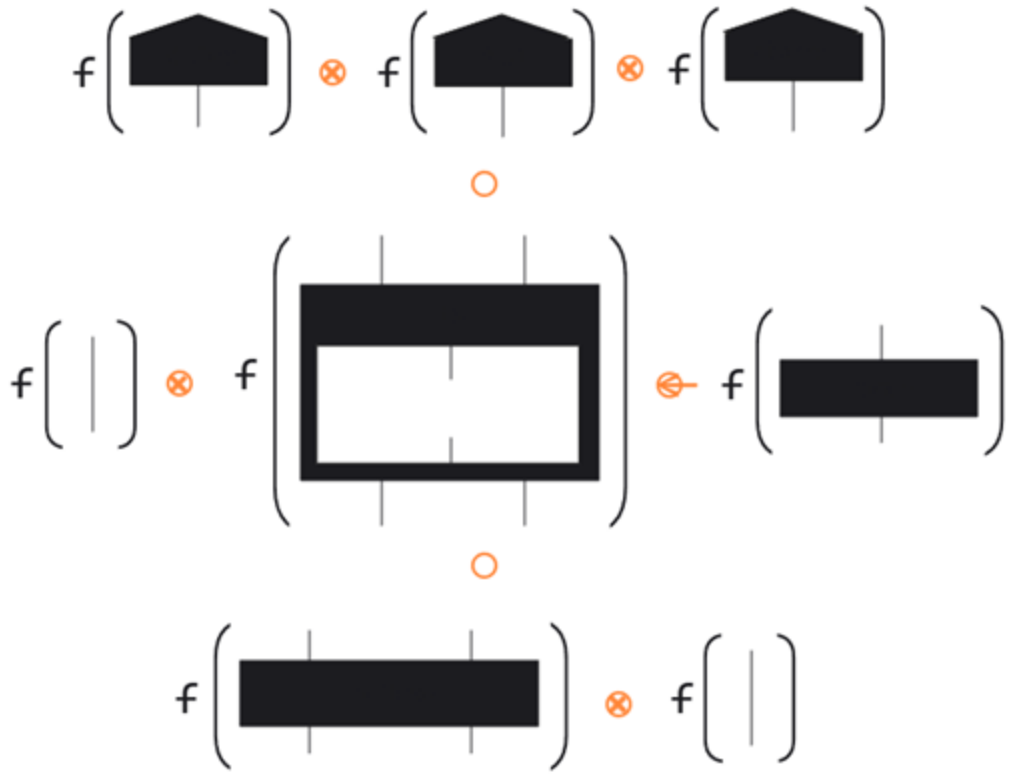

Many black boxes

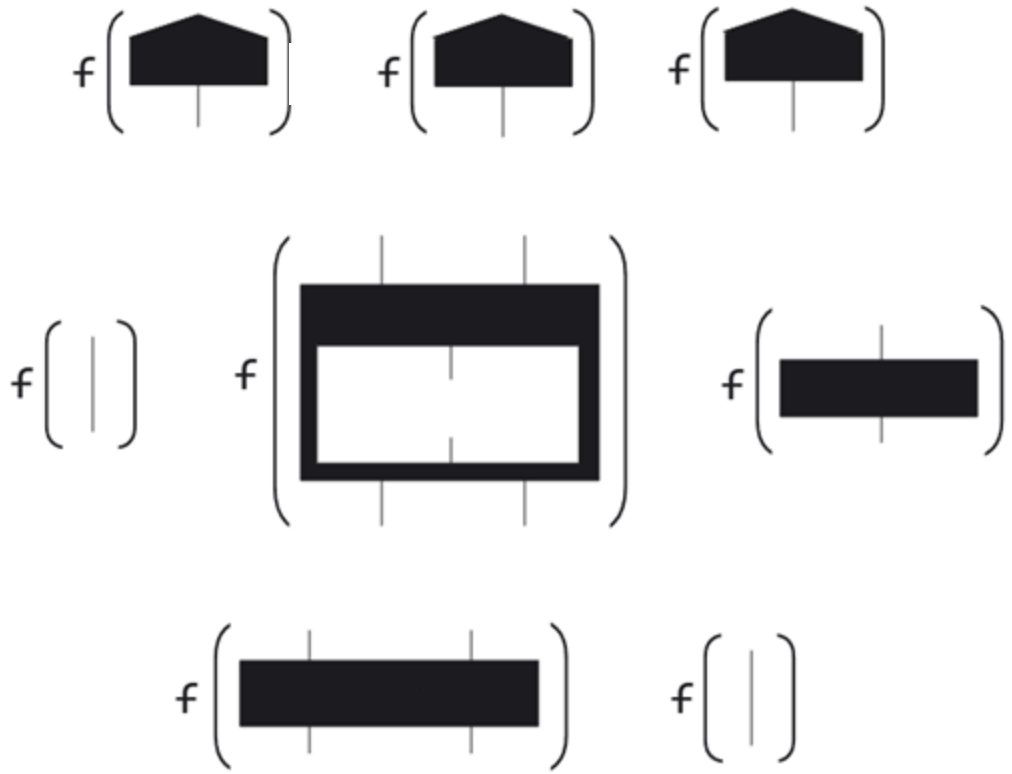

…decomposed

Many labelled black boxes

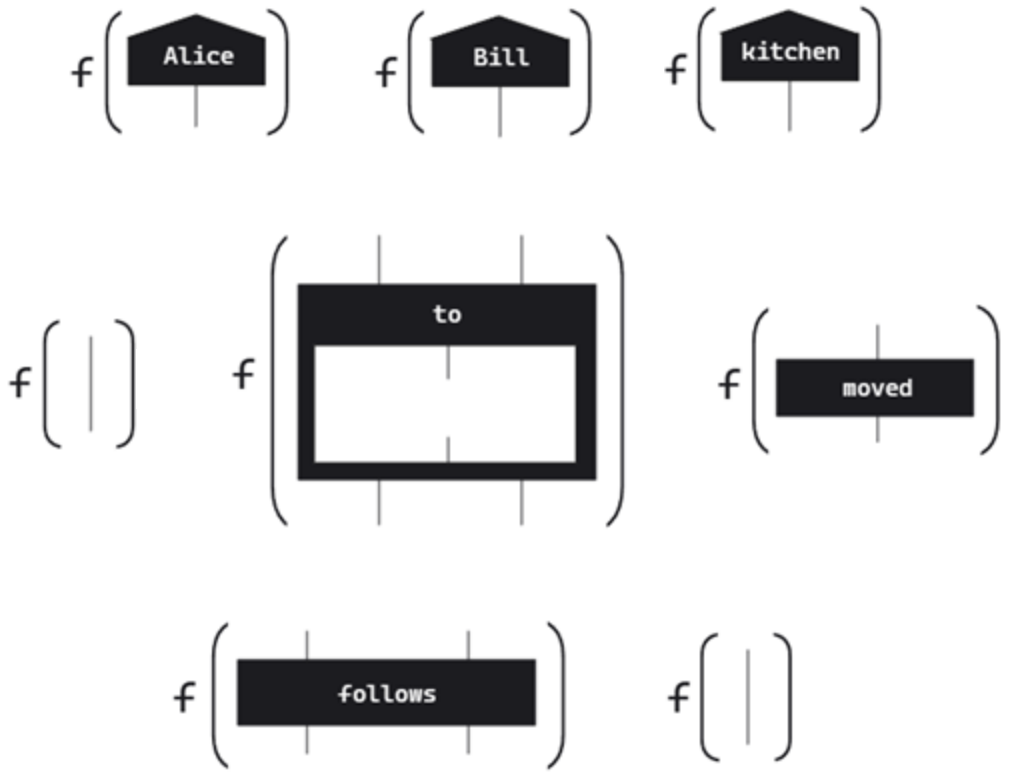

…decomposed

Many labelled black boxes

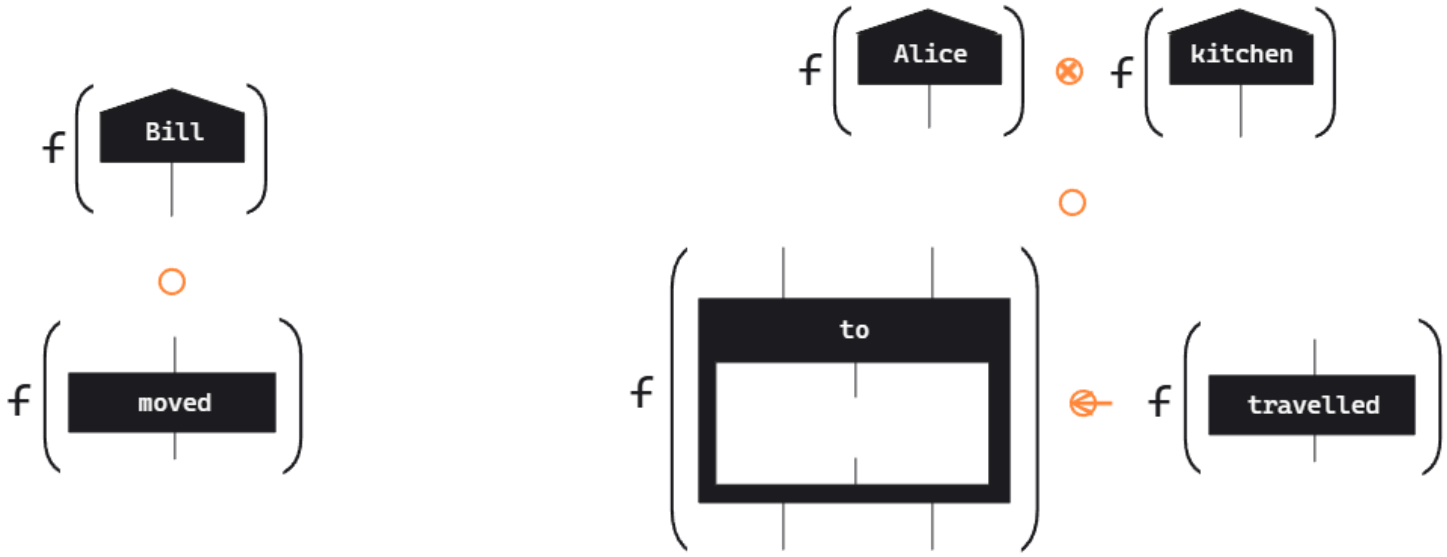

…recomposed

Productivity: bAbI 6

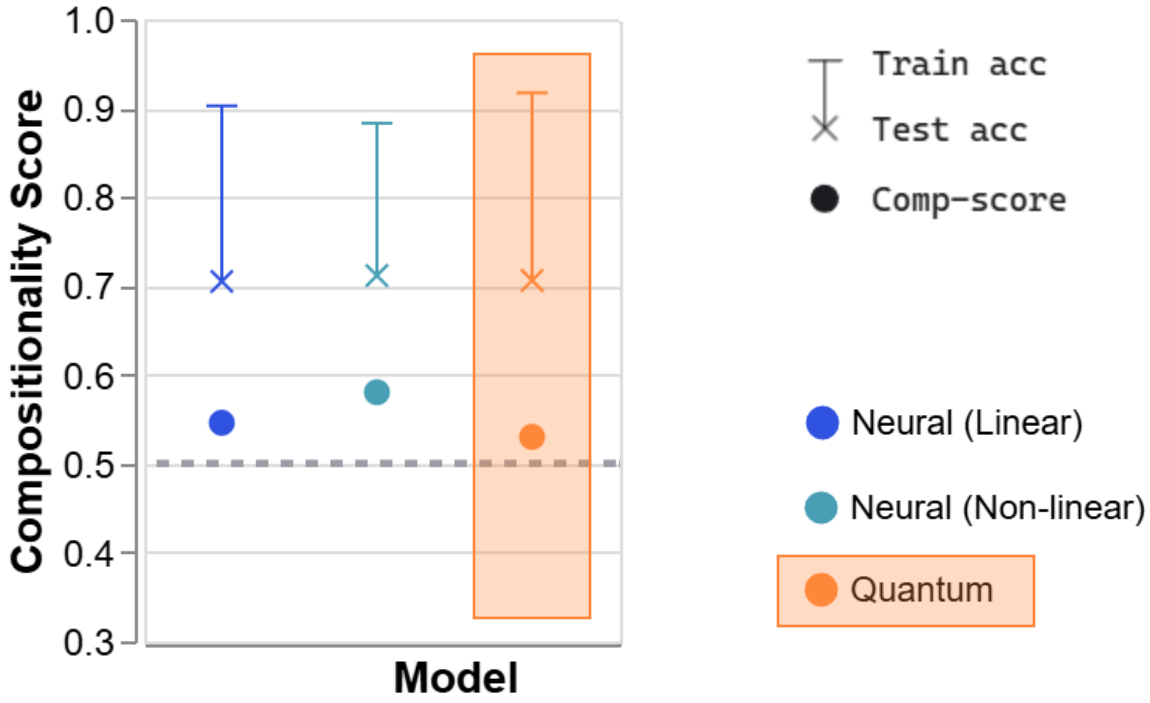

Diagram fragment

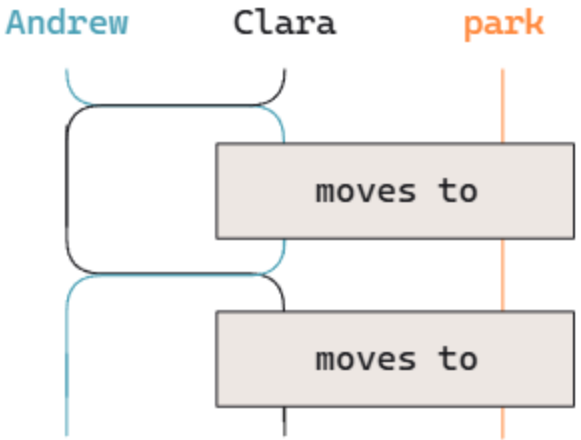

Andrew moved to the park.
Clara moved to the park.

Axioms

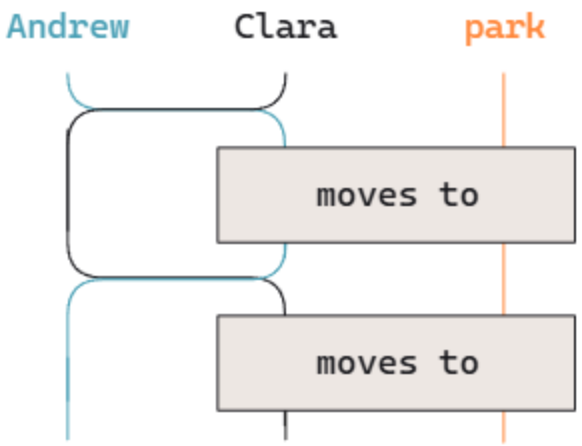

Andrew moved to the park.
Clara moved to the park.

=

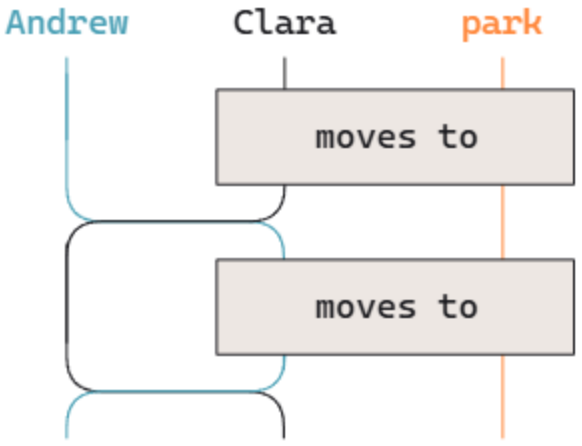

Clara moved to the park.
Andrew moved to the park.

Assertion-relative fragment

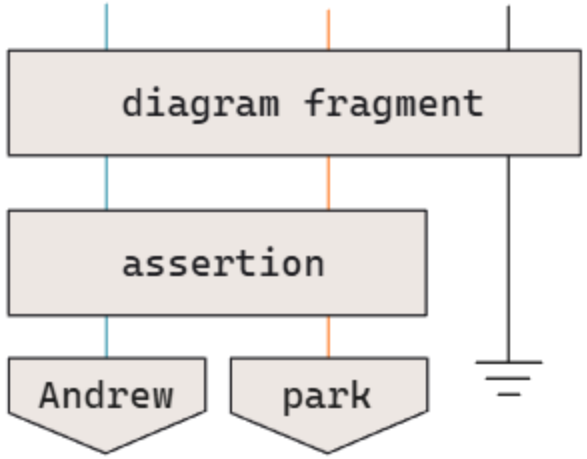

Is Andrew in the park?

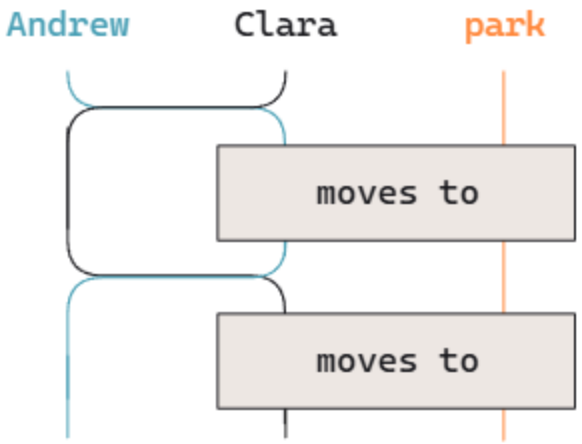

Andrew moved to the park.
Clara moved to the park.

=

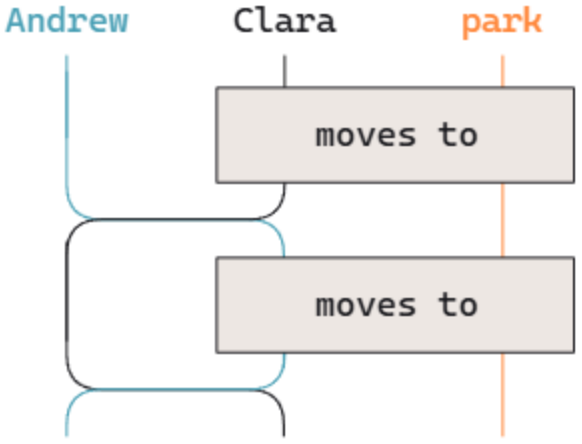

Clara moved to the park.
Andrew moved to the park.

Is Andrew in the park?

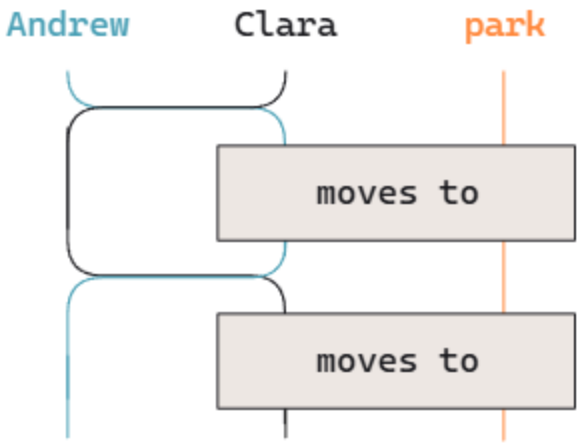

Andrew moved to the park.
Clara moved to the park.

≠

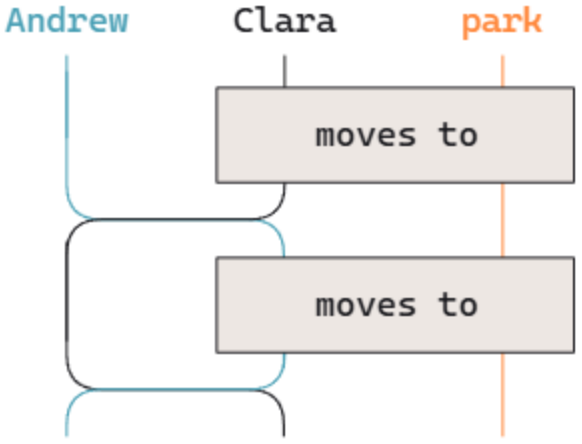

Clara moved to the park.
Andrew moved to the park.

Is ~~Andrew~~ somebody in the park?

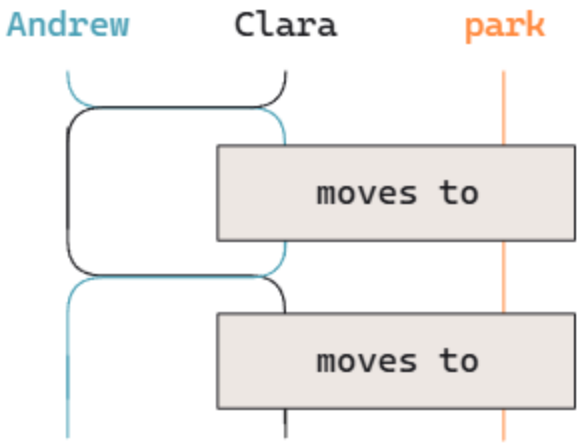

Andrew moved to the park.
Clara moved to the park.

≠

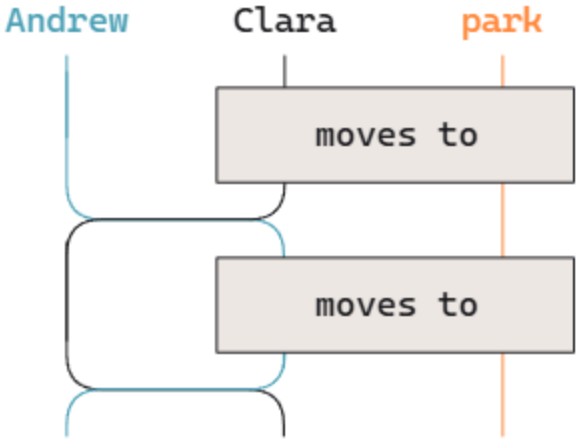

Clara moved to the park.
Andrew moved to the park.

Is ~~Andrew~~ somebody in the park?

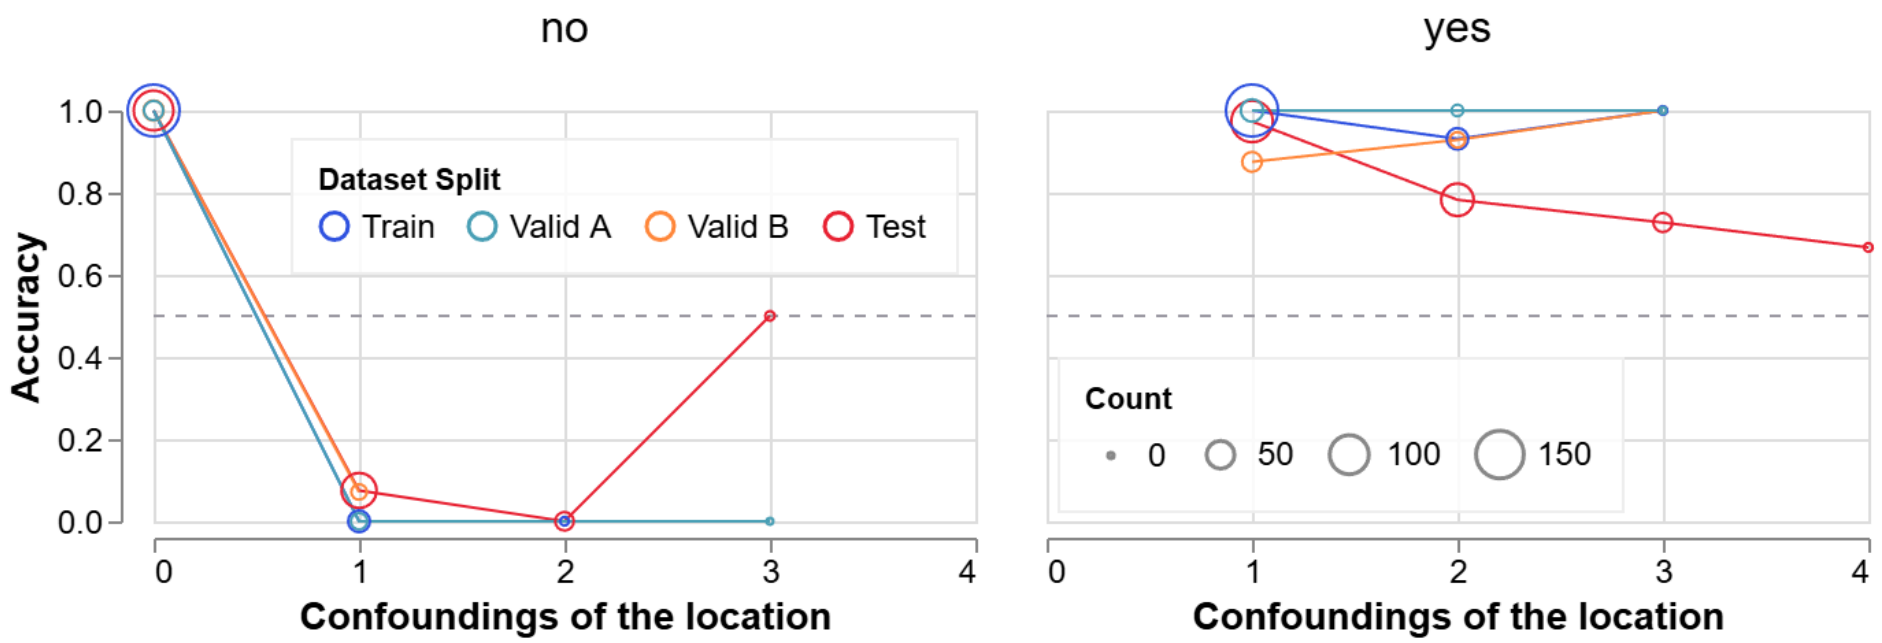

Summary

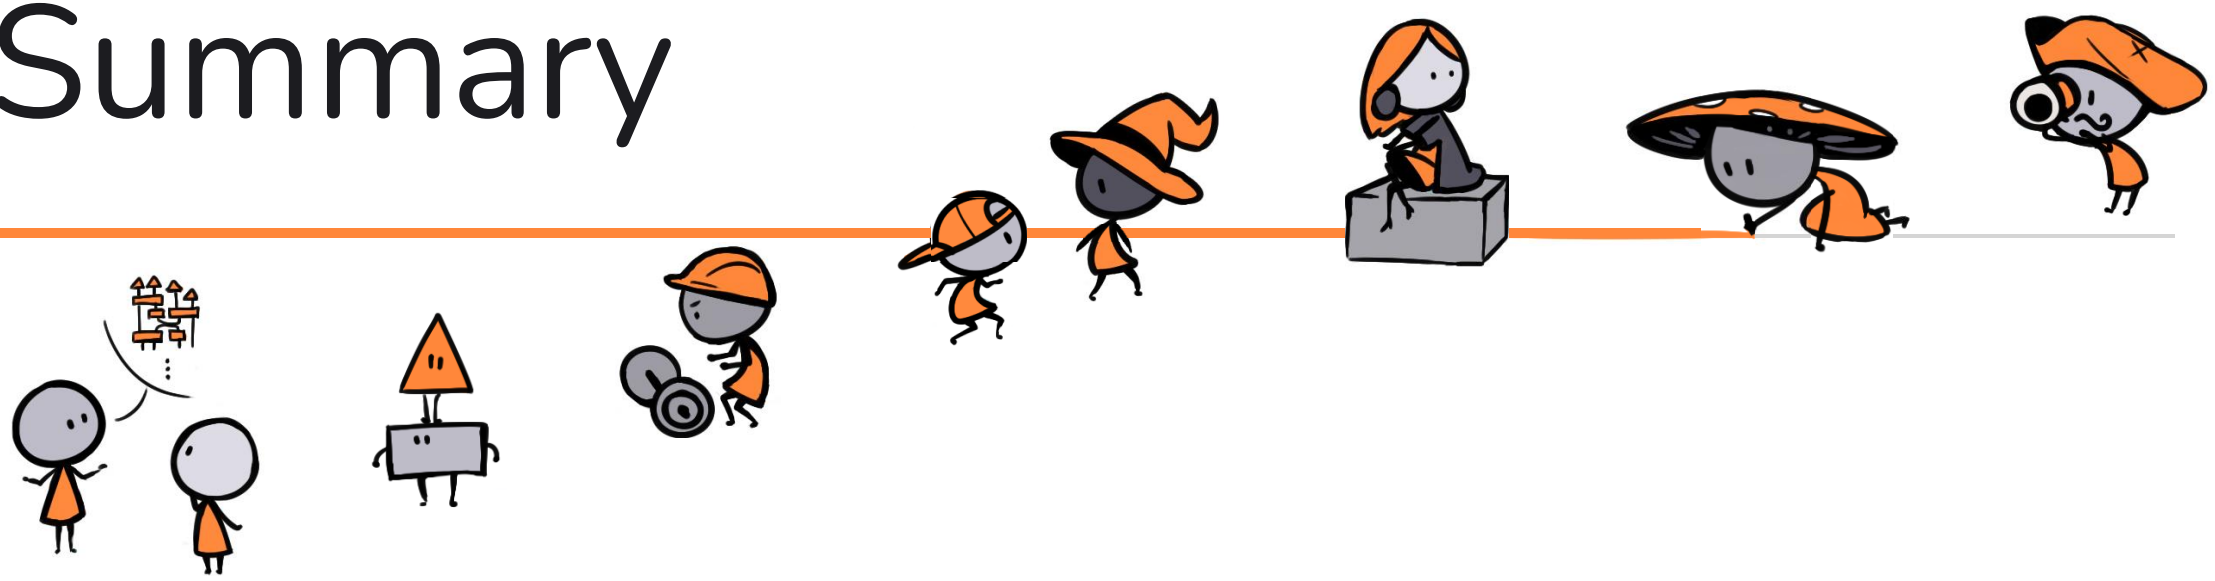

Compositionality

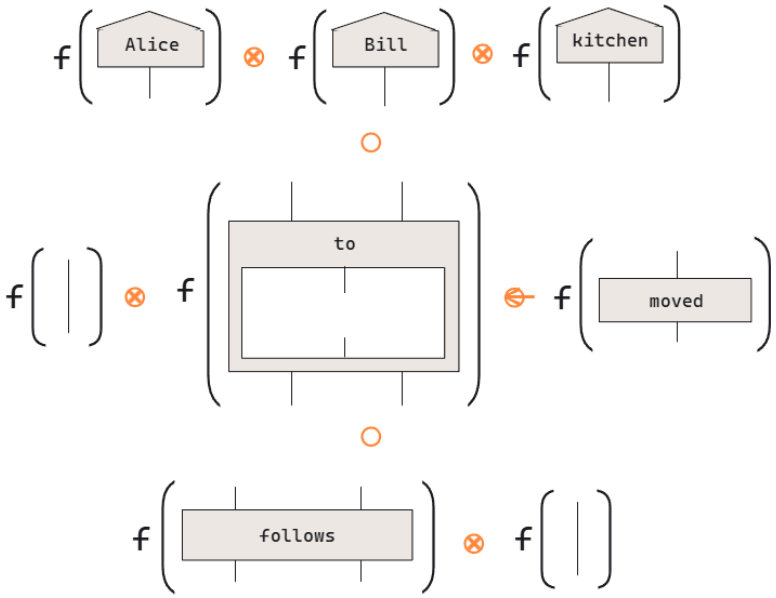

Compositionality

Task

↓

Compositional Generalisation

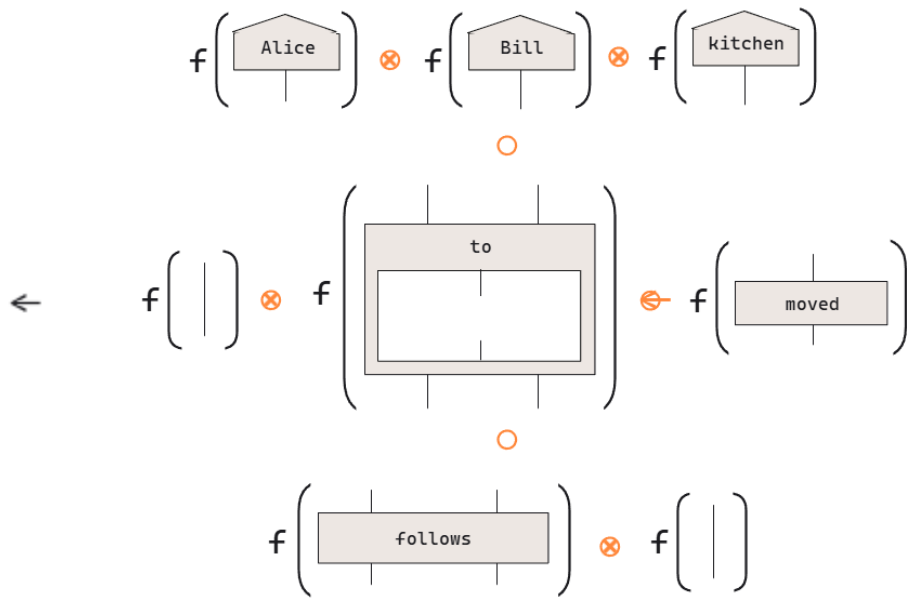

Compositionality

Task

↓

Compositional Generalisation

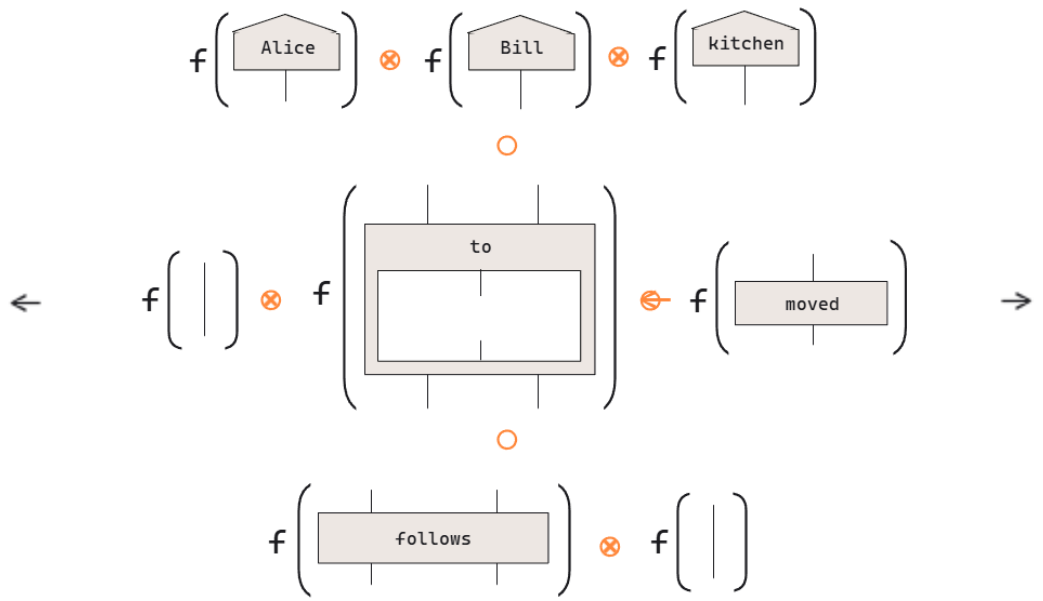

Model

↓

Compositional Interpretability

Compositionality

Task

↓

Compositional Generalisation

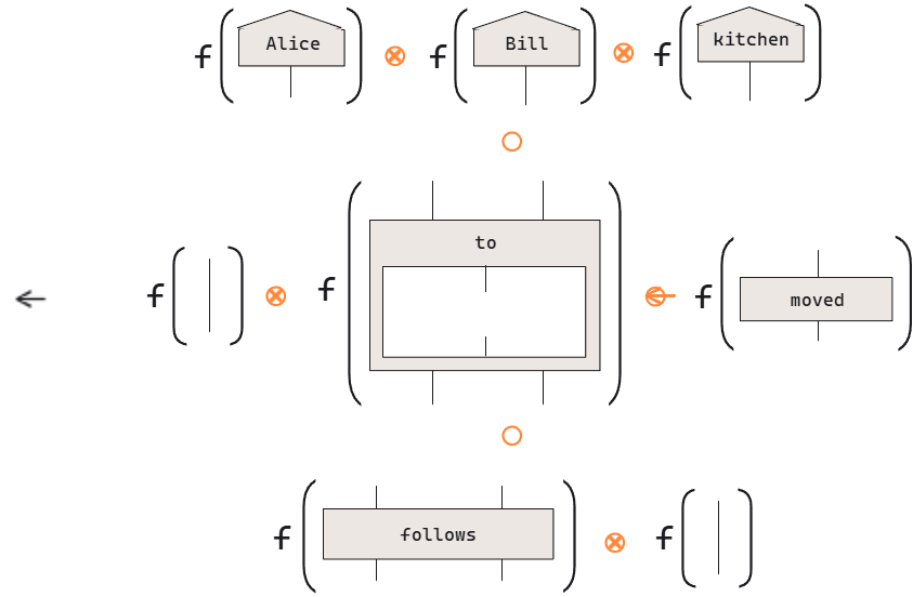

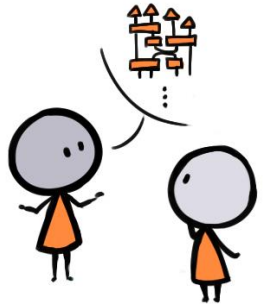

Model

↓

Compositional Interpretability

Bonus Slides

Systematicity

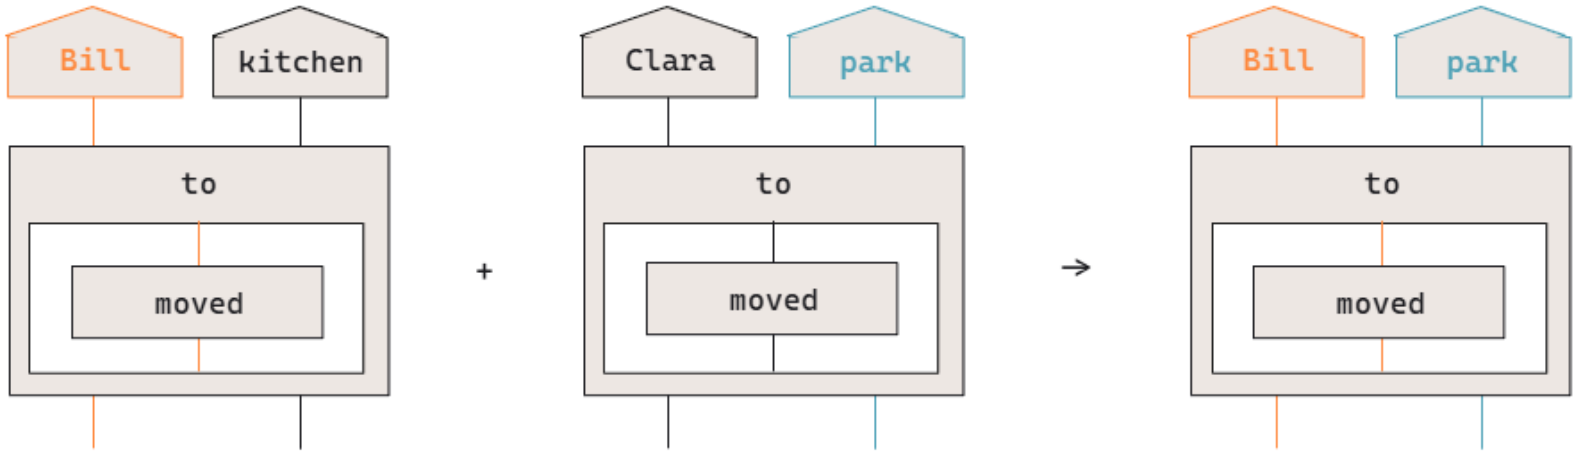

Substitutivity

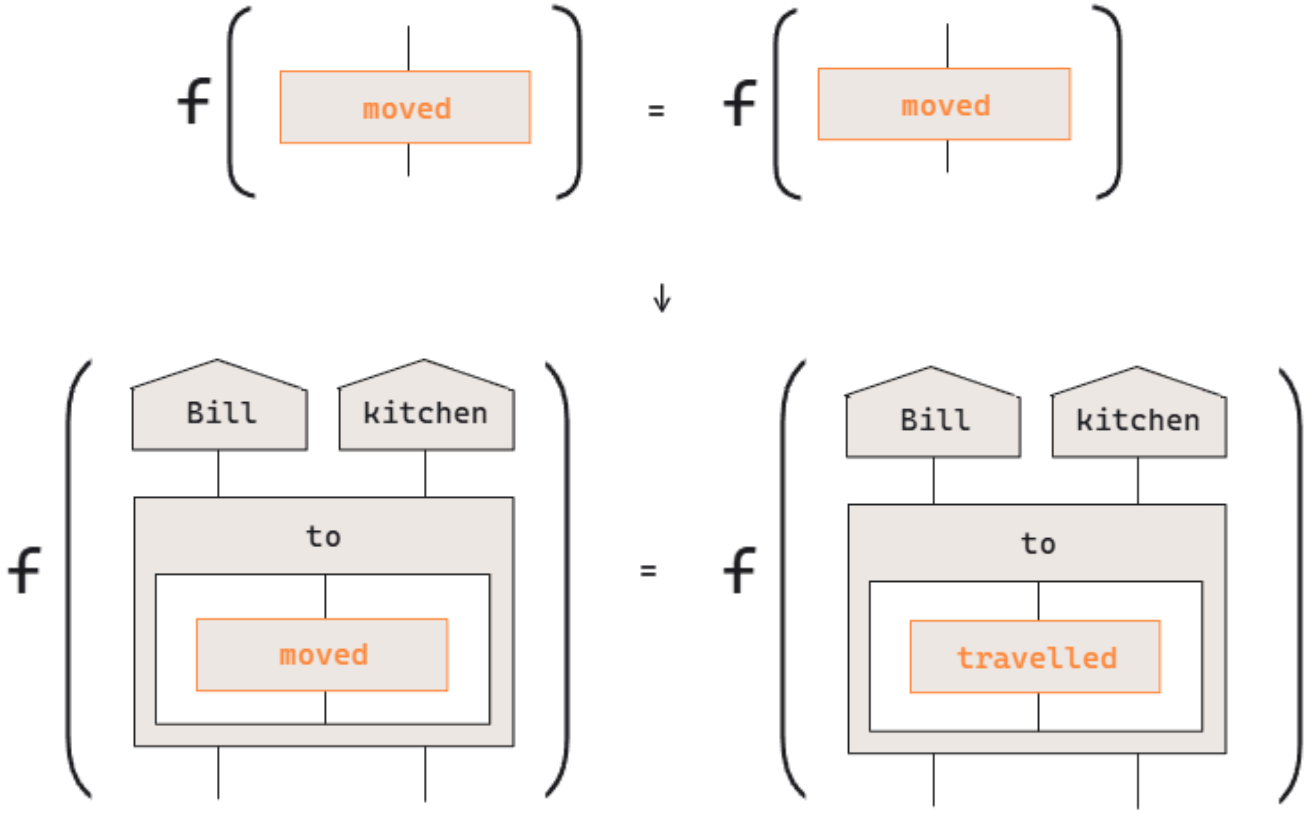

Overgeneralisation

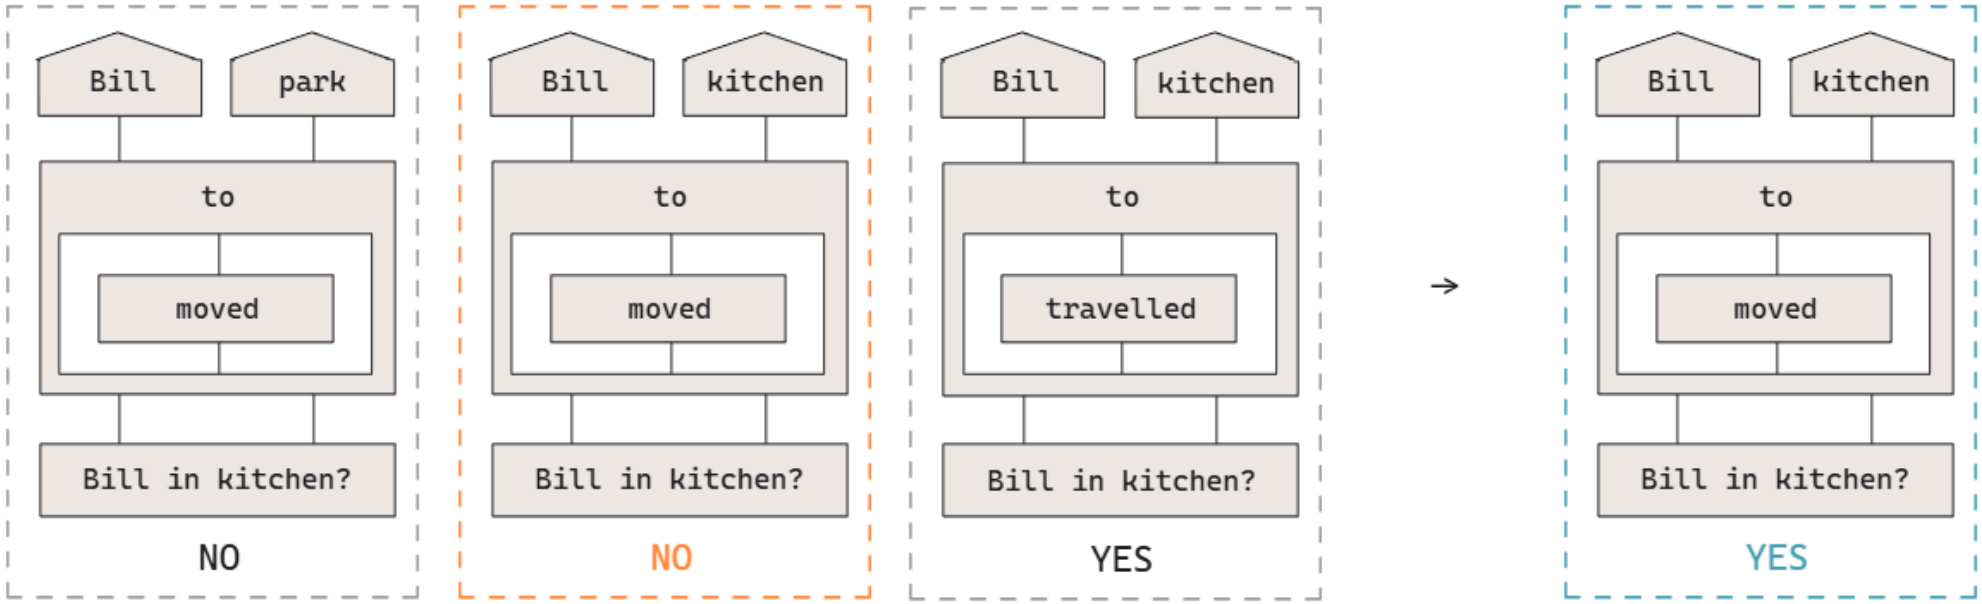

Substitutivity

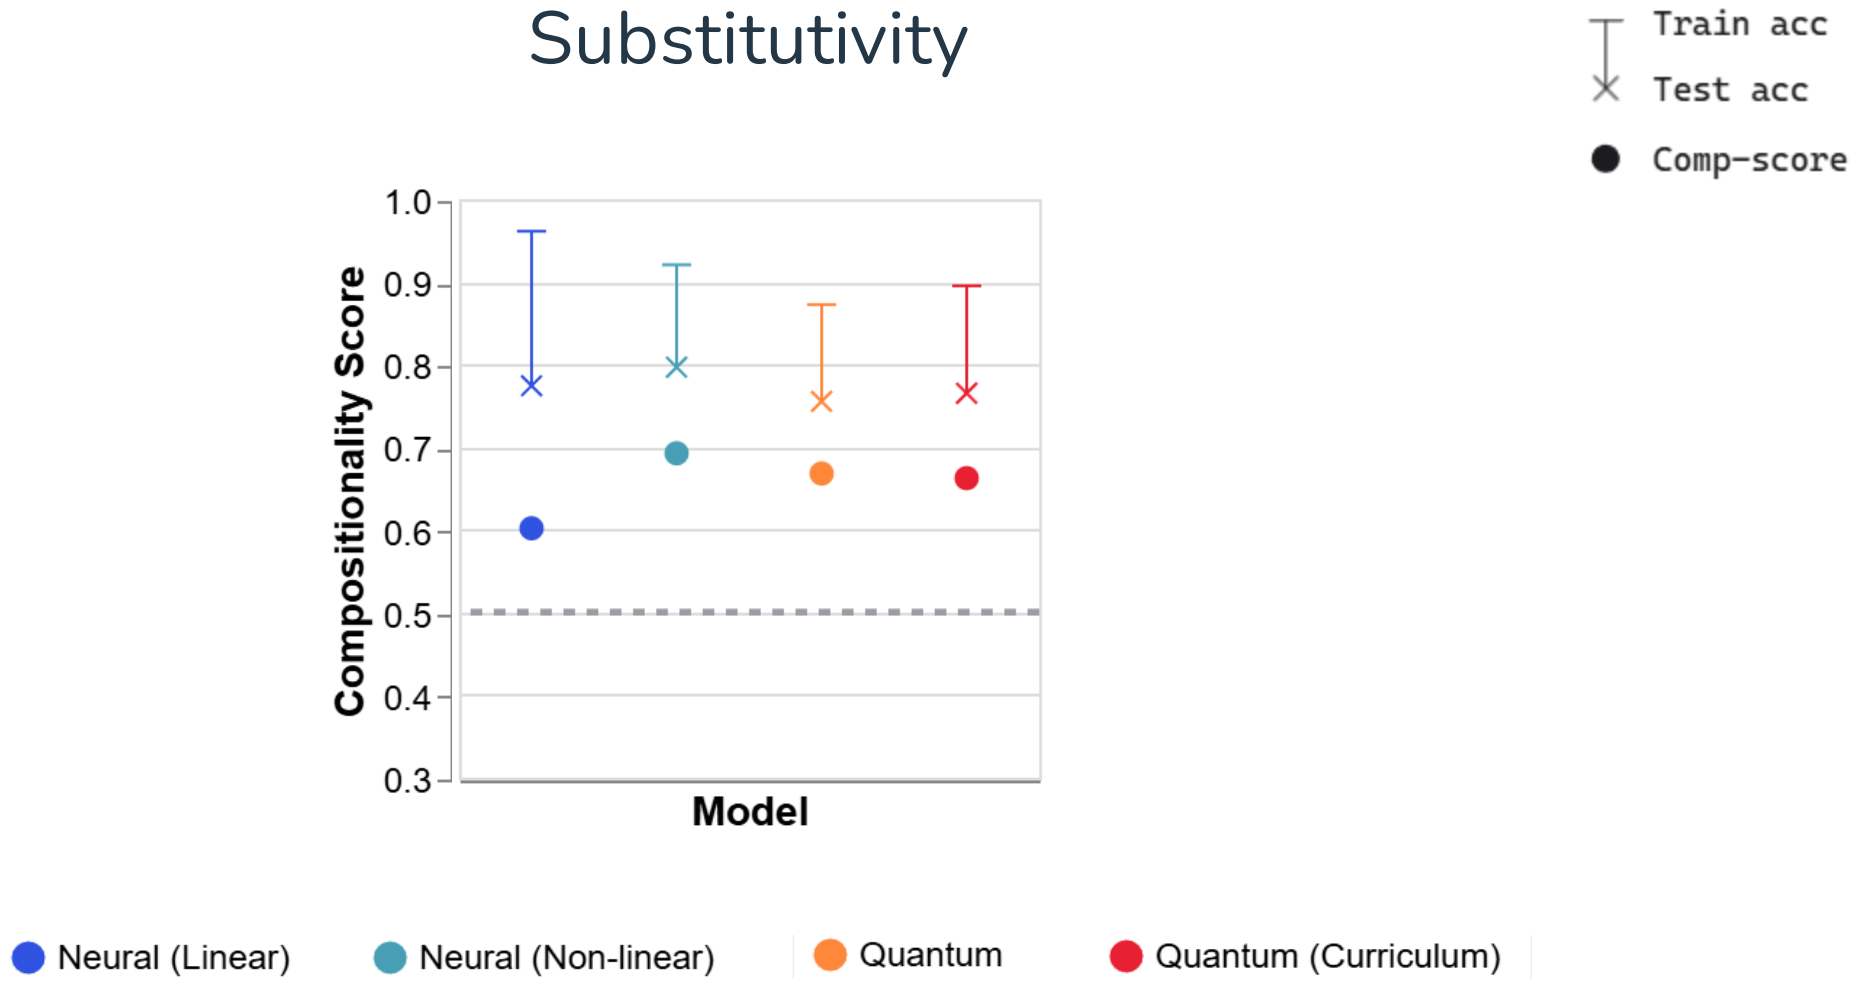

Substitutivity

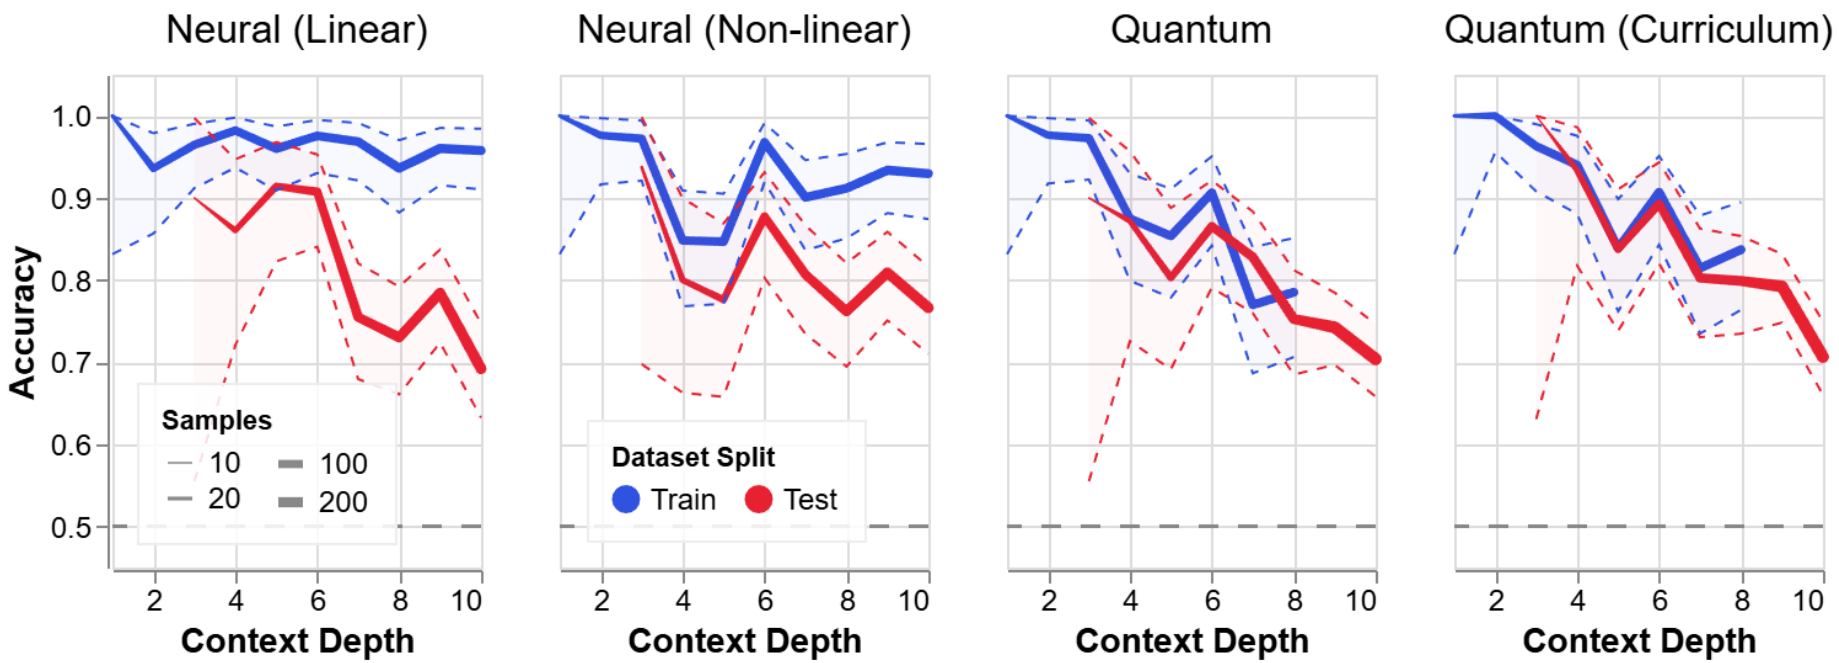

Systematicity

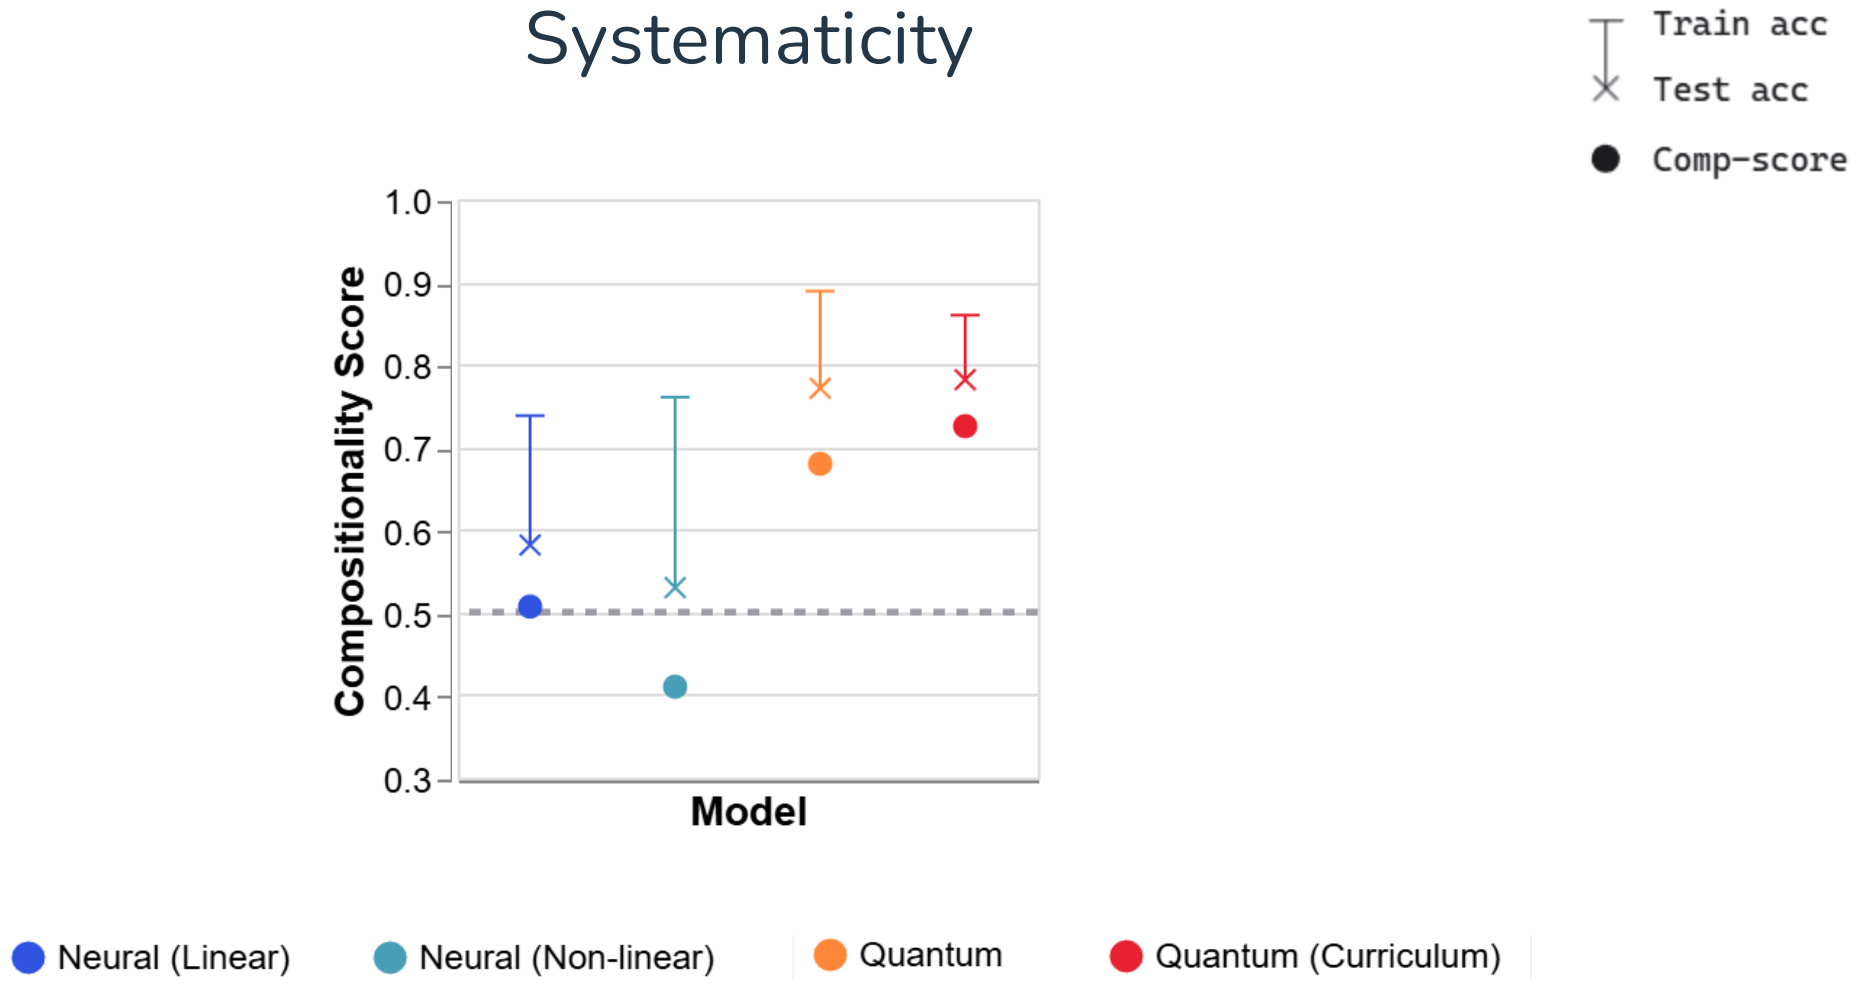

Overgeneralisation (ID)

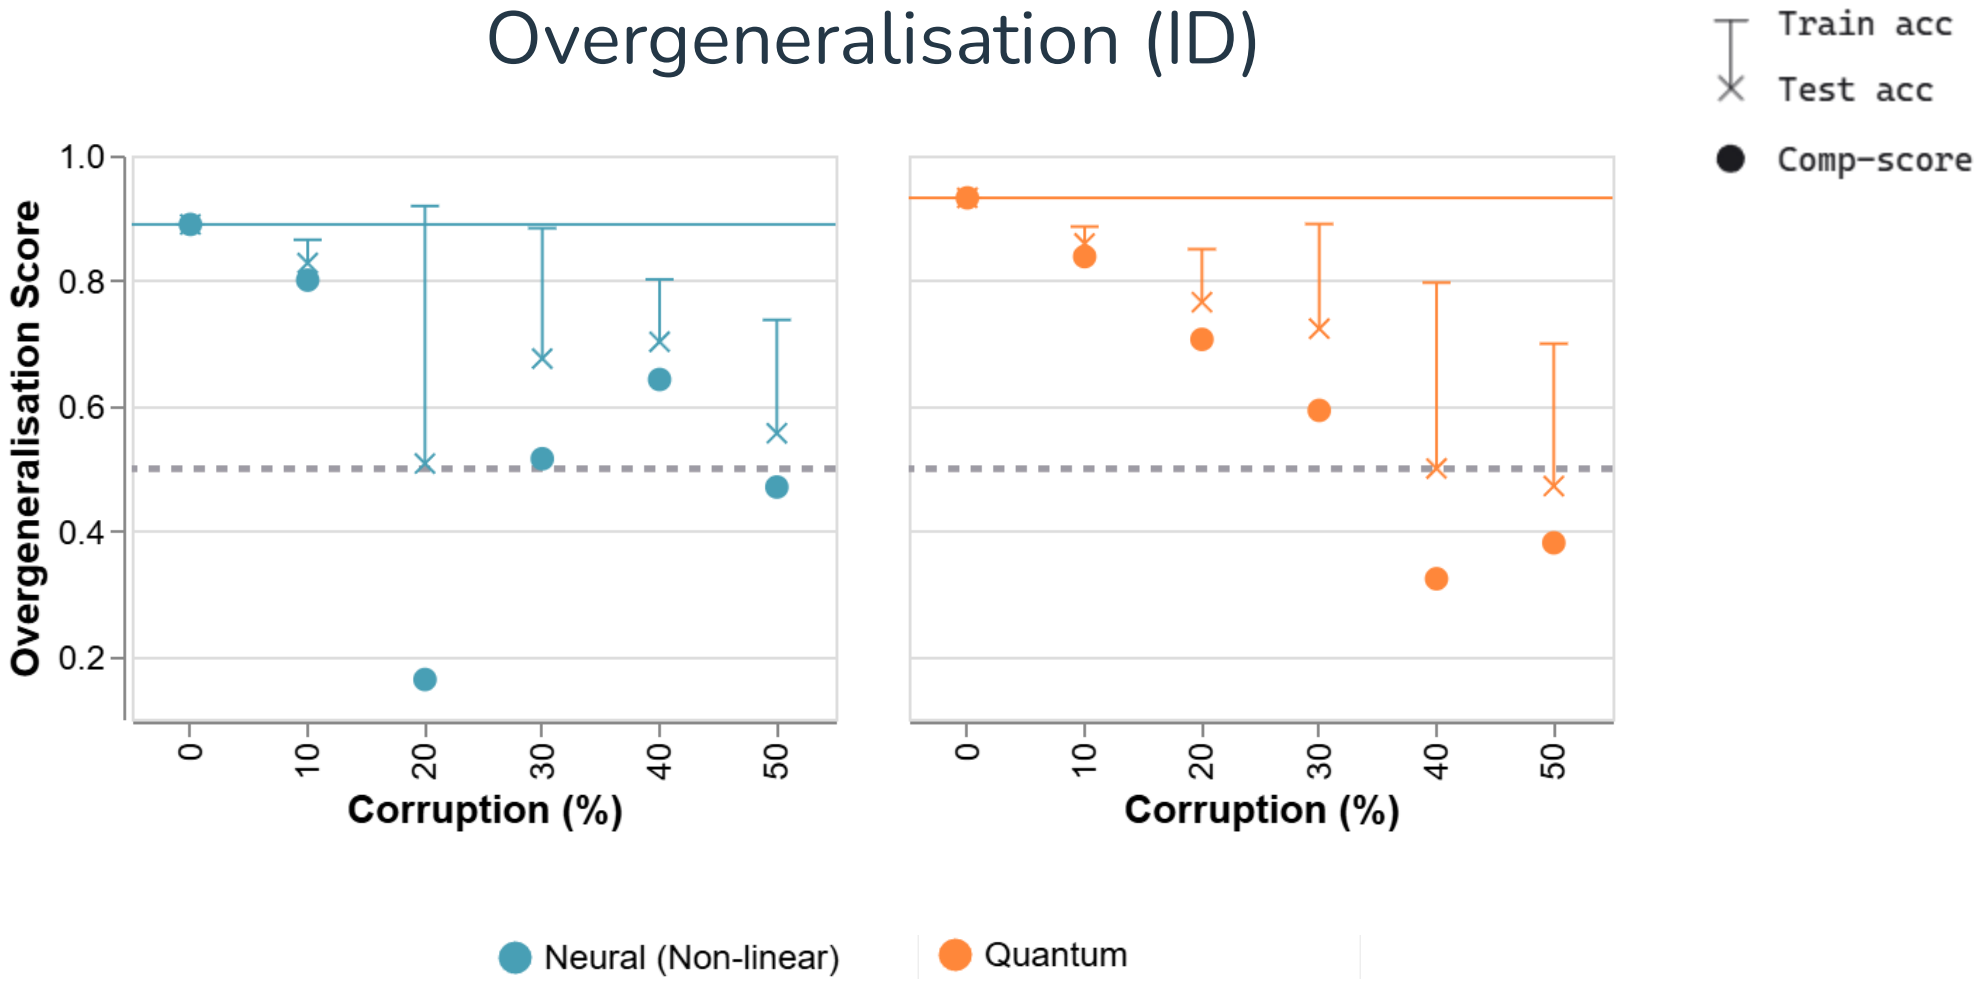

Overgeneralisation (OOD)

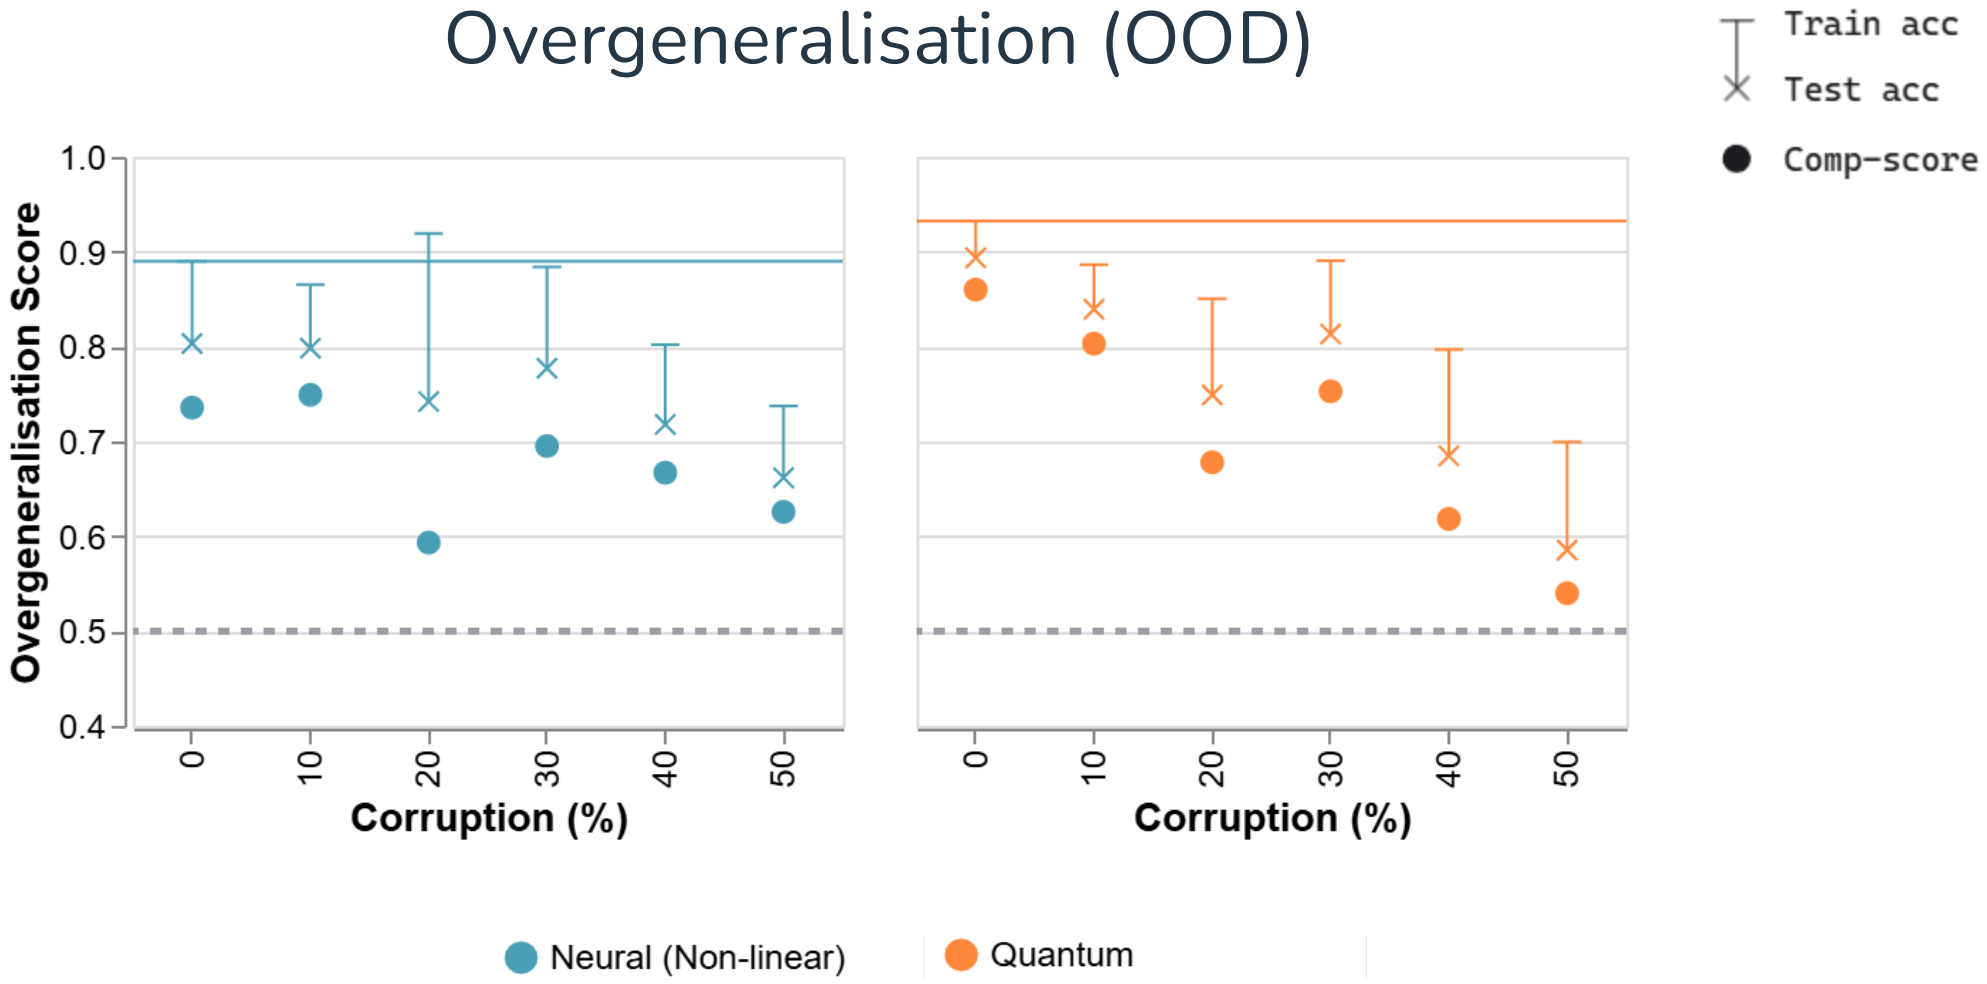

Productivity: bAbI 6

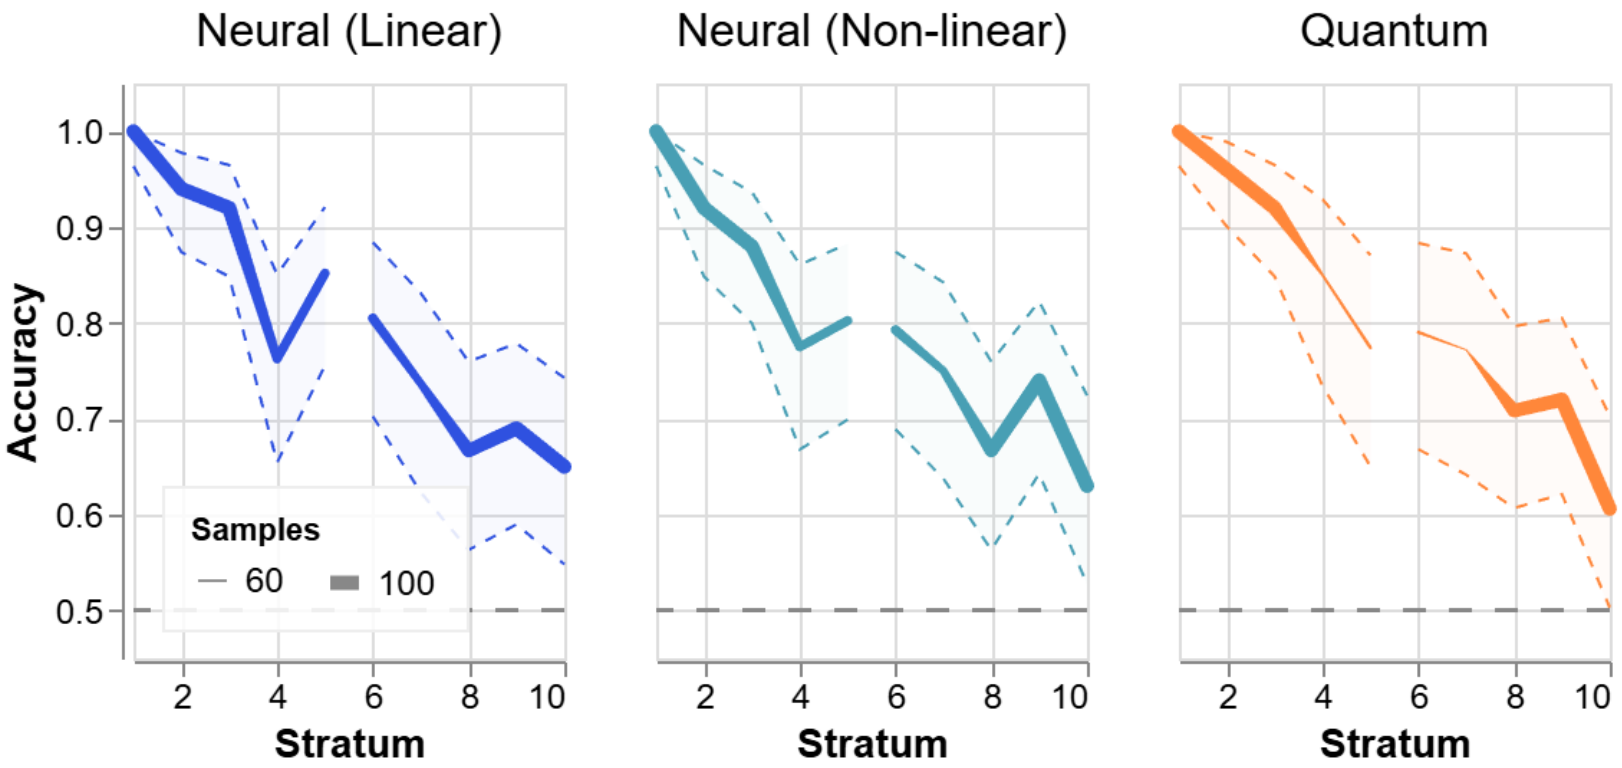

Pipeline

TKET

QUANTINUUM
H1-1

via

QUANTINUUM
NEXUS

Tensor
Network

qujax

2

directions

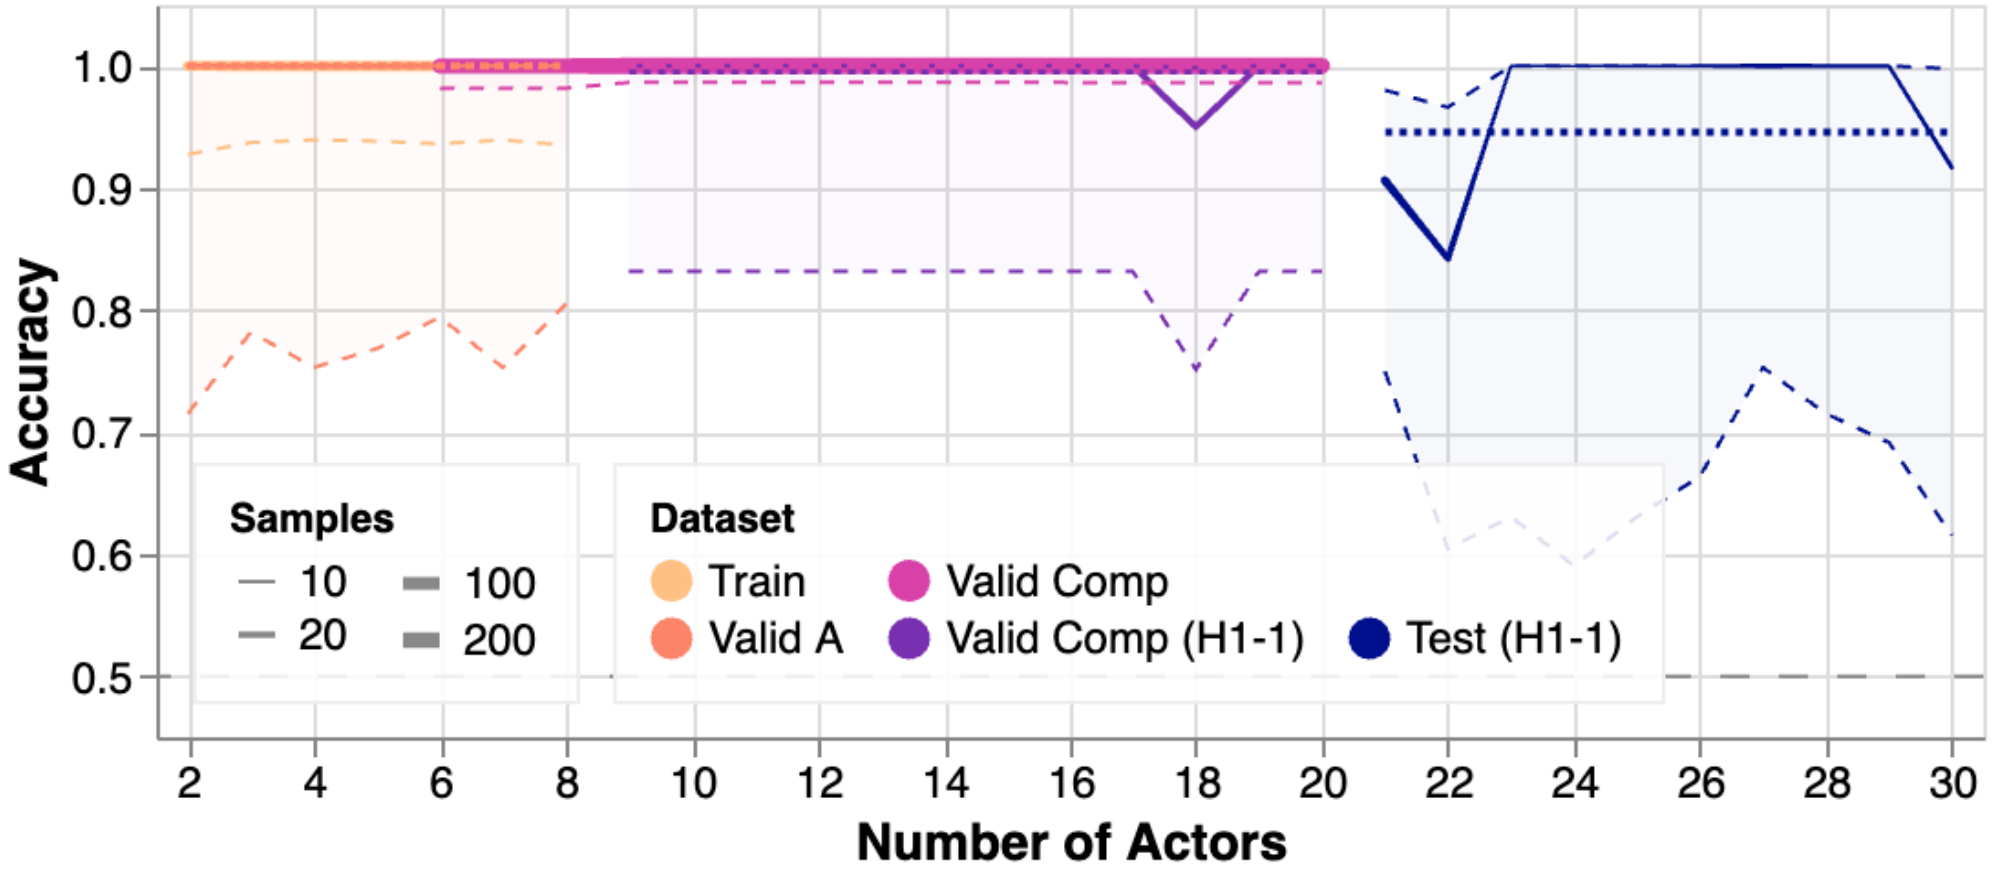

2

directions

Evaluating Test dataset

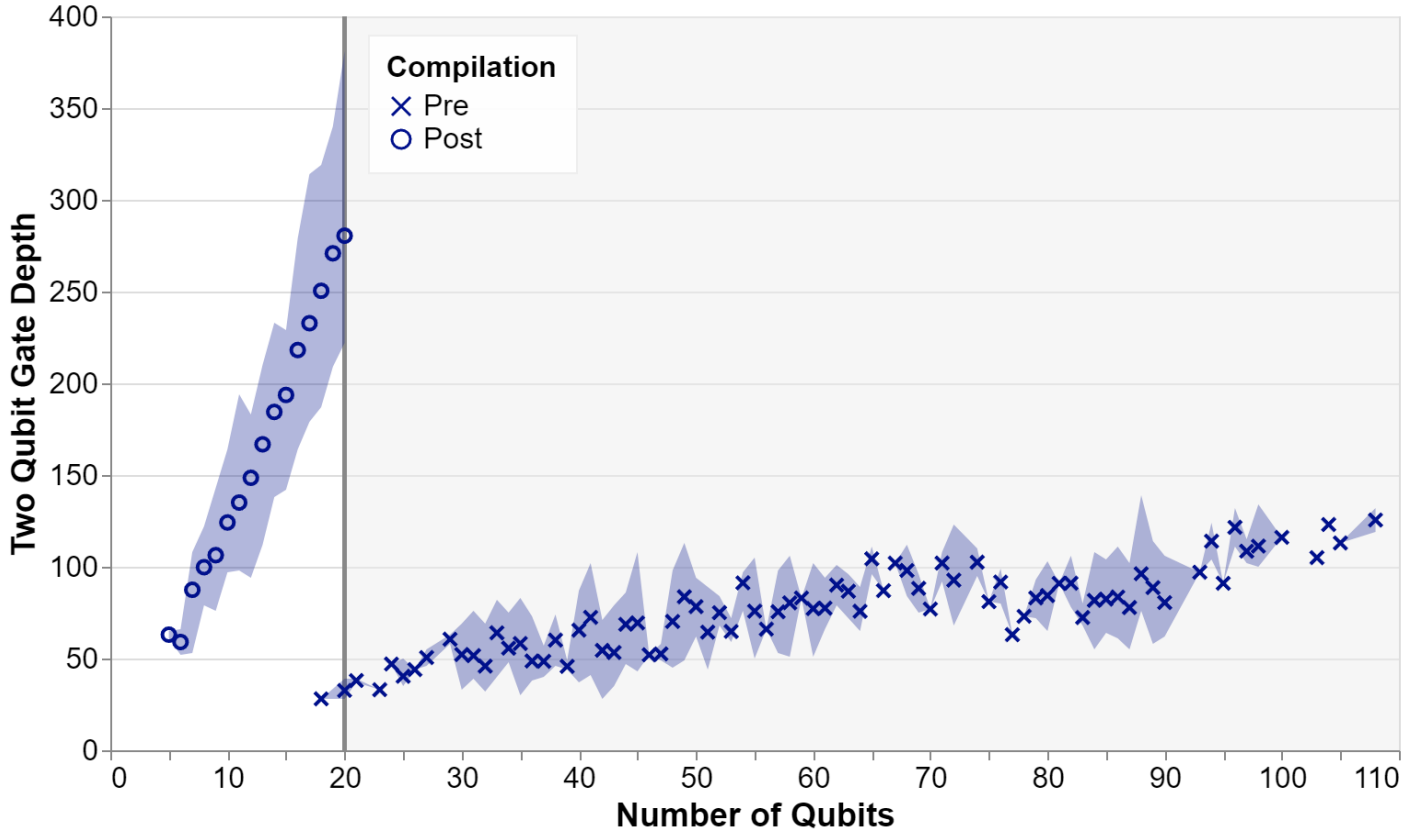

TKET

https://docs.quantinuum.com/tket/

4

directions

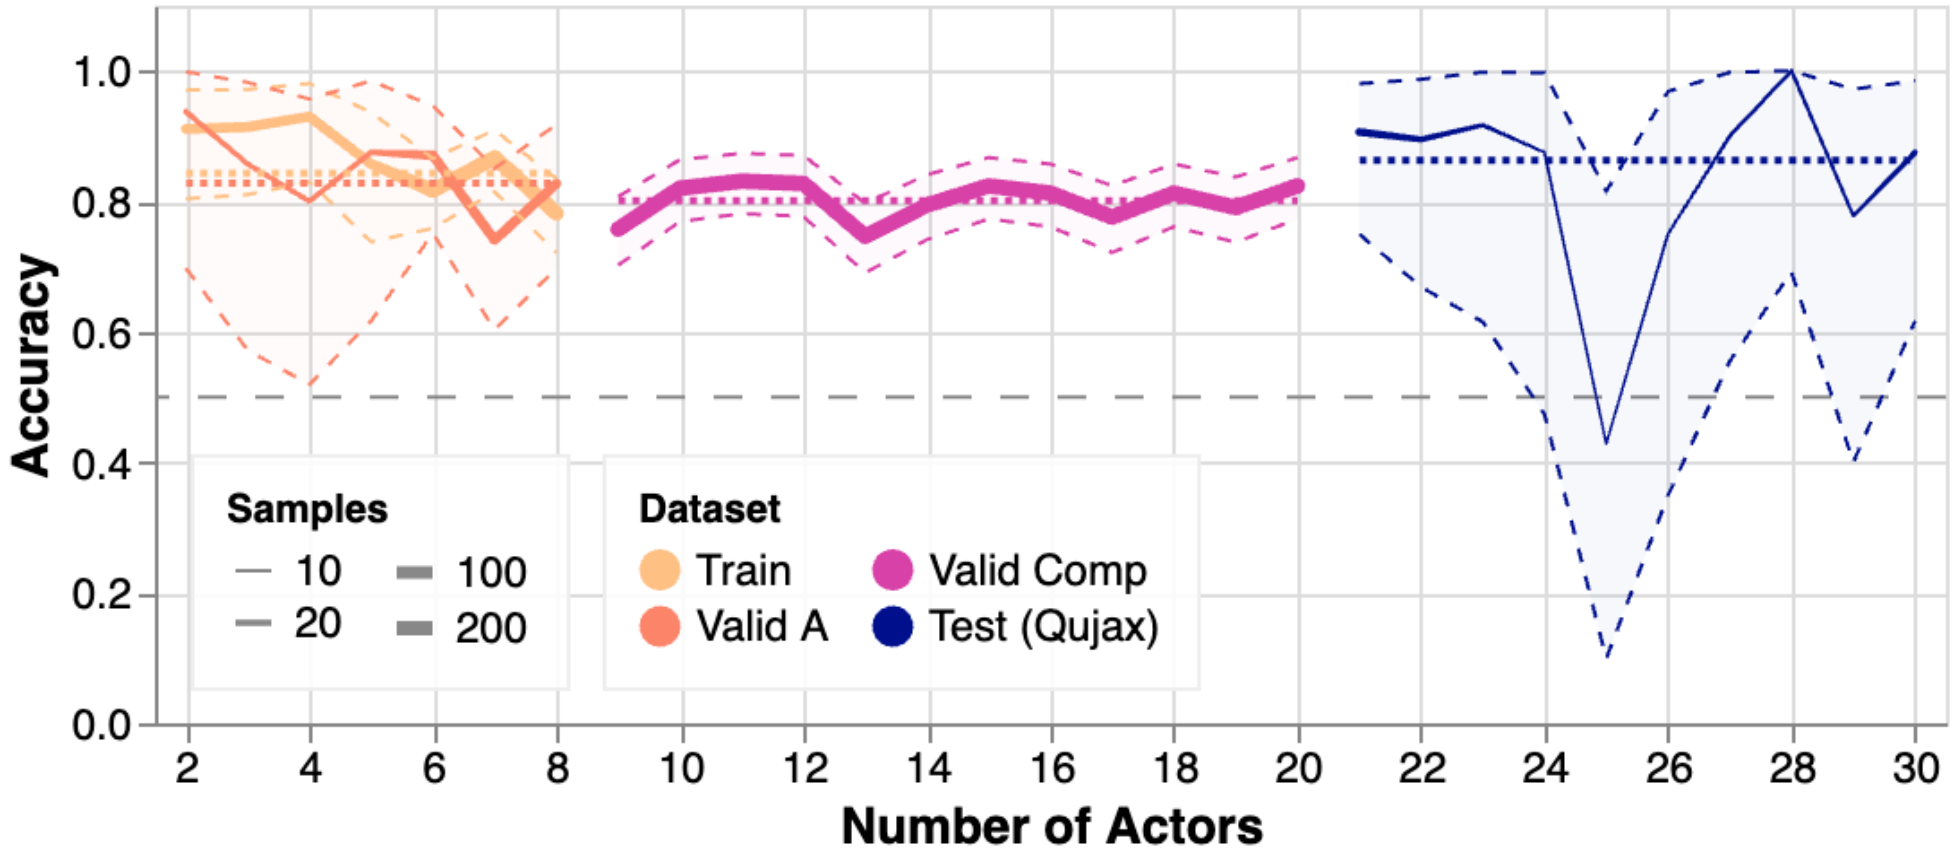

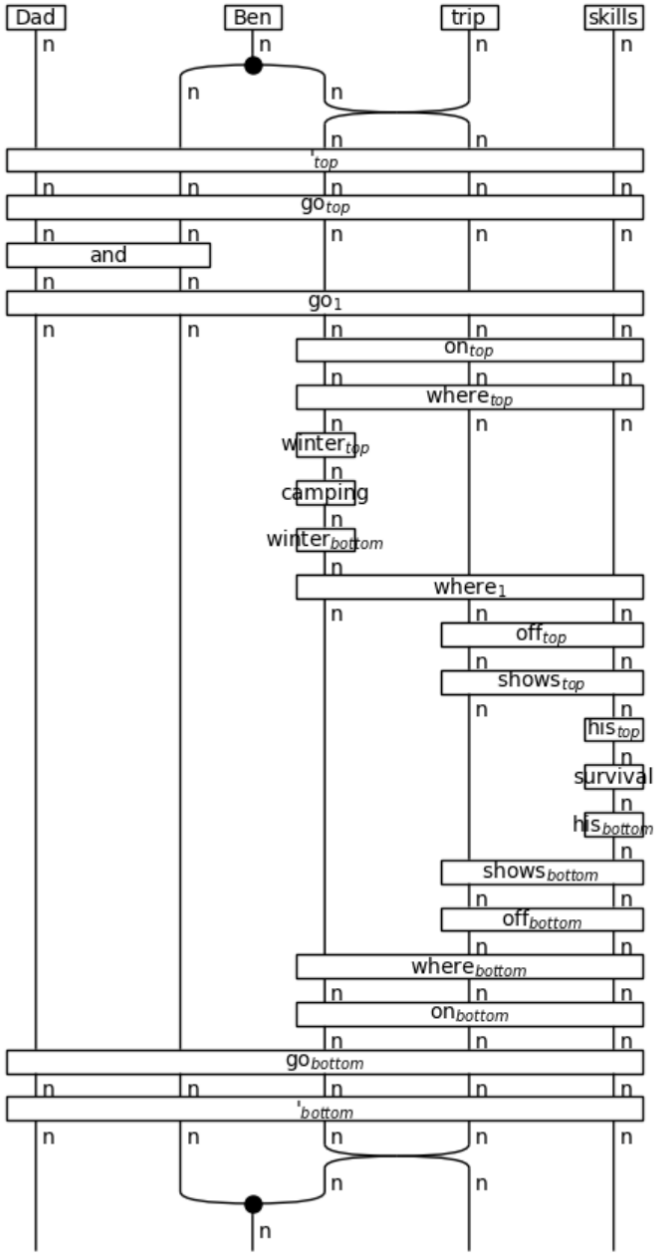

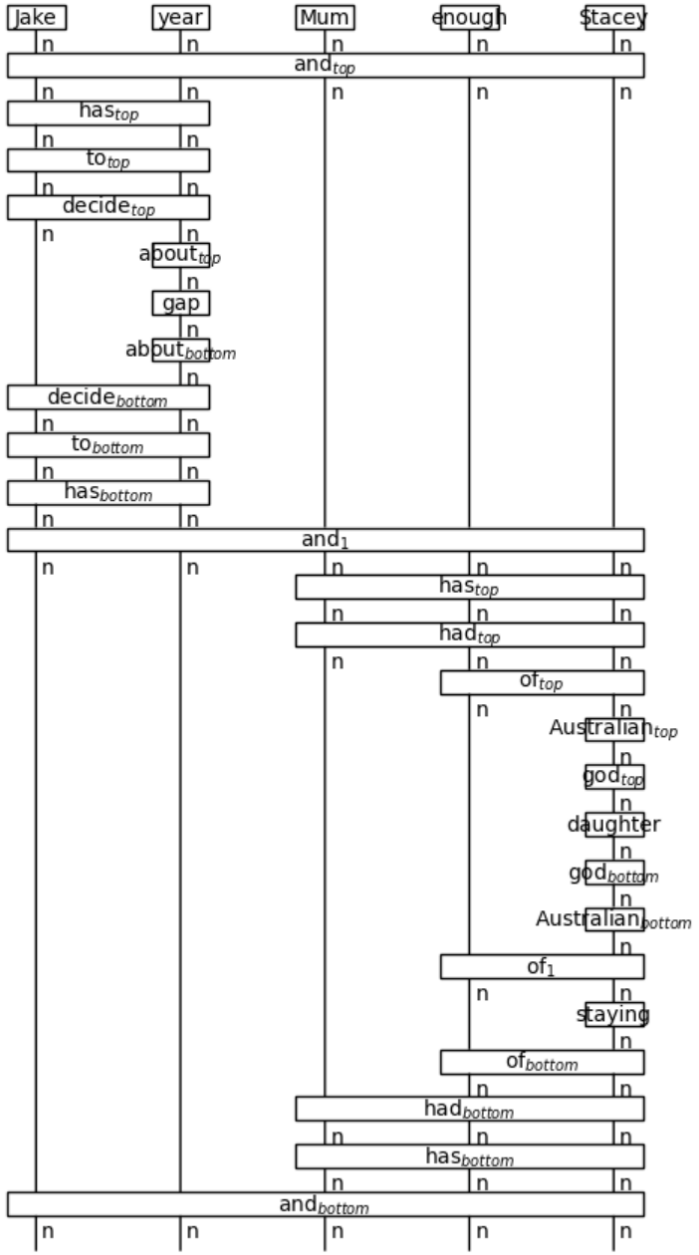

Real Data

"Dad and Ben go on a winter camping trip, where Ben shows off his survival skills. Jake has to decide about a gap year, and Mum has had enough of Australian god-daughter Stacey staying."

- BBC Synopsis Dataset