



# Are Neural Networks Compositional, and How Do We Even Know?

Dieuwke Hupkes



Yann Dubois

## Location Attention for Extrapolation to Longer Sequences

**Yann Dubois**  
University of Cambridge  
yann.dubois96@gmail.com

**Gautier Dagan**  
University of Amsterdam  
gaudier.dagan@gmail.com

**Dieuwke Hupkes\***  
ILLC, University of Amsterdam  
d.hupkes@uva.nl

**Elia Bruni\***  
Universitat Pompeu Fabra  
elia.bruni@gmail.com



Gautier Dagan



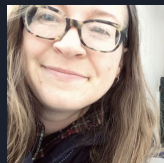
Mario Giulianelli

## Under the Hood: Using Diagnostic Classifiers to Investigate and Improve how Language Models Track Agreement Information

**Mario Giulianelli** **Jack Harding** **Florian Mohnert**  
University of Amsterdam University of Amsterdam University of Amsterdam  
{mario.giulianelli, jack.harding, florian.mohnert}@student.uva.nl

**Dieuwke Hupkes**  
ILLC, University of Amsterdam  
d.hupkes@uva.nl

**Willem Zuidema**  
ILLC, University of Amsterdam  
w.h.zuidema@uva.nl



Kate McCurdy

## Generalising to German Plural Noun Classes, from the Perspective of a Recurrent Neural Network

**Verna Dankers\***, **Anna Langedijk\***, **Kate McCurdy\***,  
**Adina Williams\***, **Dieuwke Hupkes\***



Koustuv Sinha

## Masked Language Modeling and the Distributional Hypothesis: Order Word Matters Pre-training for Little

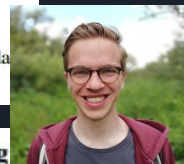
**Koustuv Sinha††** **Robin Jia†** **Dieuwke Hupkes†** **Joelle Pineau††**

**Adina Williams†** **Douwe Kiela†**  
† Facebook AI Research; † McGill University / Mila - Quebec AI  
{koustuvs, adinawilliams, dkiela}@fb.com

## Do Language Models Understand *Anything*? On the Ability of LSTMs to Understand Negative Polarity Items

**Jaap Jumelet**  
University of Amsterdam  
jaap.jumelet@student.uva.nl

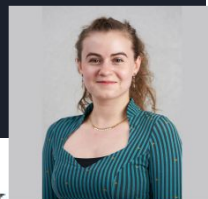
**Dieuwke Hupkes**  
ILLC, University of Amsterdam  
d.hupkes@uva.nl



Jaap Jumelet

## Language Models Use Monotonicity to Assess NPI Licensing

**Jaap Jumelet†** **Milica Denić†** **Jakub Szymanik†**  
**Dieuwke Hupkes\*** **Shane Steinert-Threlkeld†**



Verna Dankers

## Compositionality Decomposed: How do Neural Networks Generalise?

## The Paradox of the Compositionality of Natural Language: A Neural Machine Translation Case Study

**Verna Dankers**  
University of Edinburgh  
vernadankers@gmail.com

**Elia Bruni**  
University of Osnabrück  
elia.bruni@gmail.com

**Dieuwke Hupkes**  
Facebook AI Research  
dieuwkehupkes@fb.com

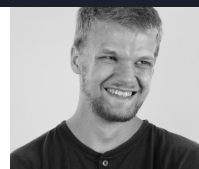
## Language Modelling as a Multi-Task Problem

**Lucas Weber**  
DTCL, University Pompeu Fabra  
lucas.weber@upf.edu


**Jaap Jumelet**  
ILLC, University of Amsterdam  
j.w.d.jumelet@uva.nl

**Elia Bruni**  
IKW, University of Osnabrück  
elia.bruni@gmail.com

**Dieuwke Hupkes**  
Facebook AI Research  
dieuwkehupkes@fb.com

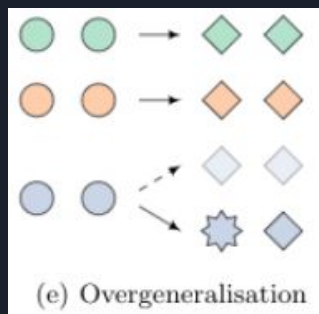
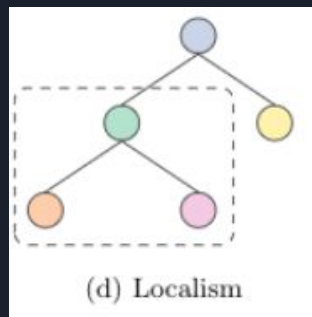
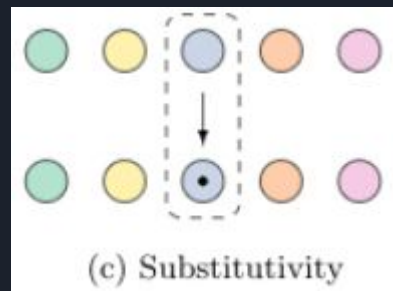
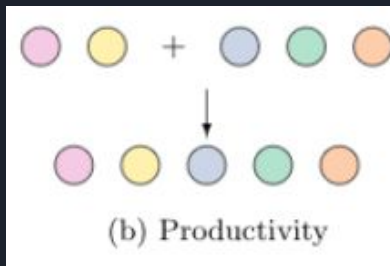
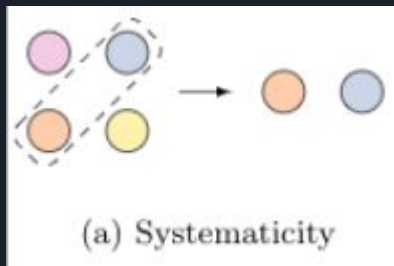


Lucas Weber

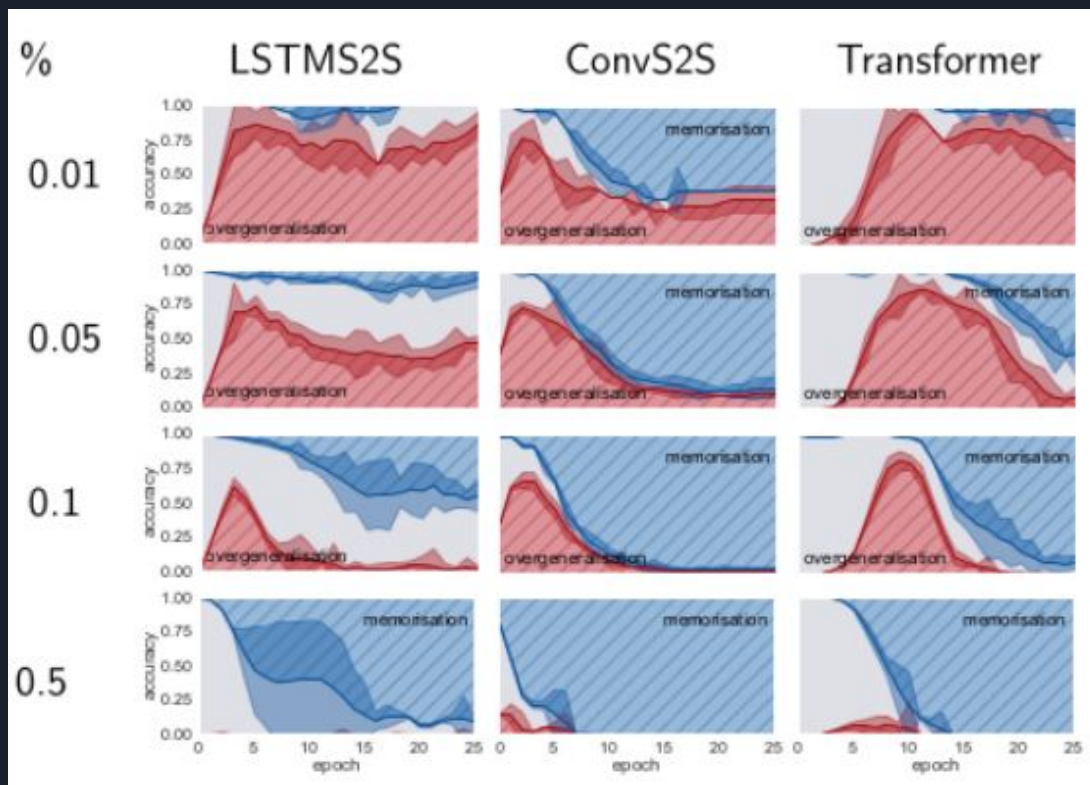


*In translation between other domains (logical languages, programming languages), compositional translation is almost trivially a sound approach, as the expressions of such languages are completely unambiguous and the compositional system according to which the meaning can be derived is known. Natural language has neither of these qualities, which does not only complicate the construction of a compositional translation system, but also renders the existence of such a system uncertain*

# How do you evaluate compositionality?



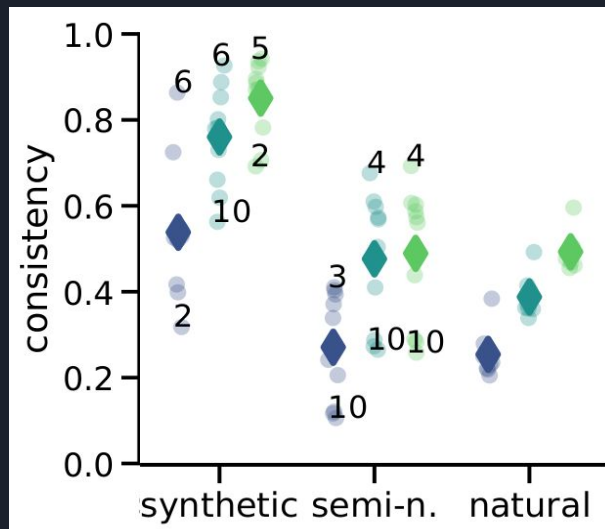
# Overgeneralisation profiles



# Compositionality in natural language

The father avoids the director, and the boys criticise the leader .

The lawyers think that the boys cry, and the boys criticise the leader .




training size

small

medium

full

S1 → S3



# An example of our substitutivity test

The mother that travels by (aeroplane|airplane) wants to facilitate the development of innovative practices in education and training at tertiary level, and their transfer, including from one participating country to others.

De moeder die per vliegtuig reist wil de ontwikkeling van innovatieve praktijken op de gebied van onderwijs en opleiding op tertiair niveau, en de overdracht daarvan, ook van de ene deelnemende land naar de andere, vergemakkelijken.

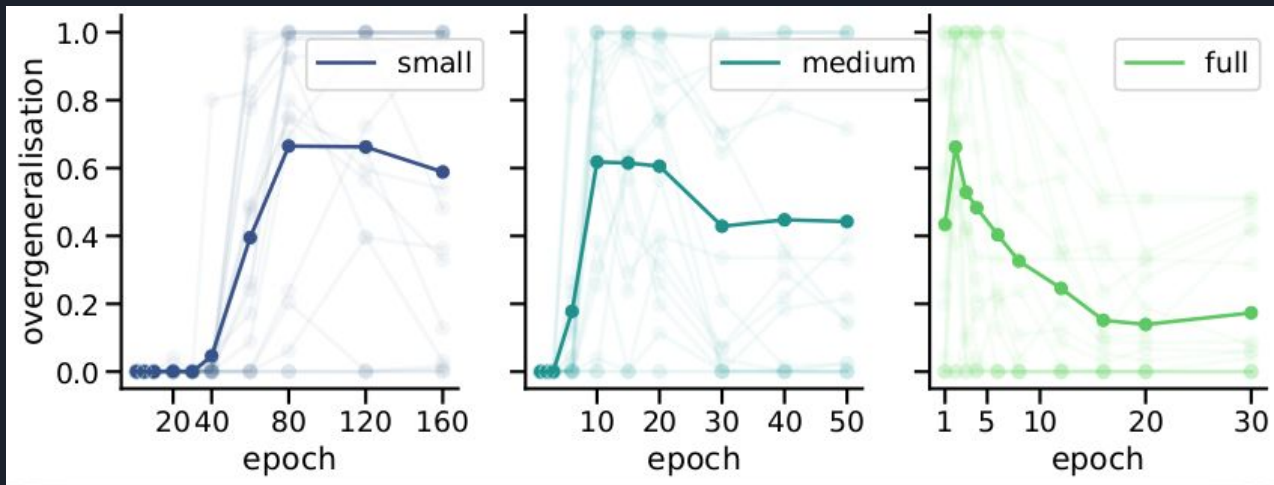
De moeder die per vliegtuig reist wil de ontwikkeling van innovatieve praktijken op de gebied van onderwijs en opleiding op tertiair niveau, en hun overdracht, waaronder van de ene deelnemende land naar de andere, vergemakkelijken.

# Overgeneralisation

You guys are out of your minds!

Jullie zijn niet goed wijs.

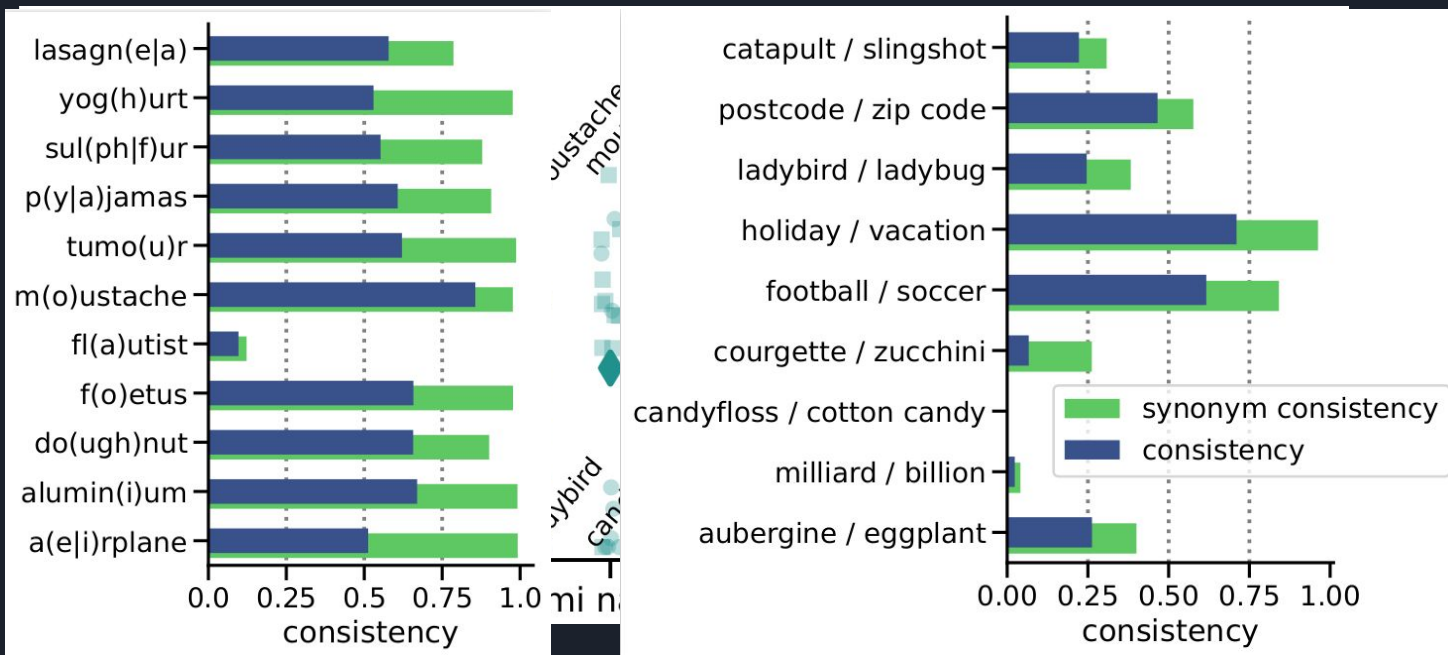
Jullie zijn uit je hoofd.



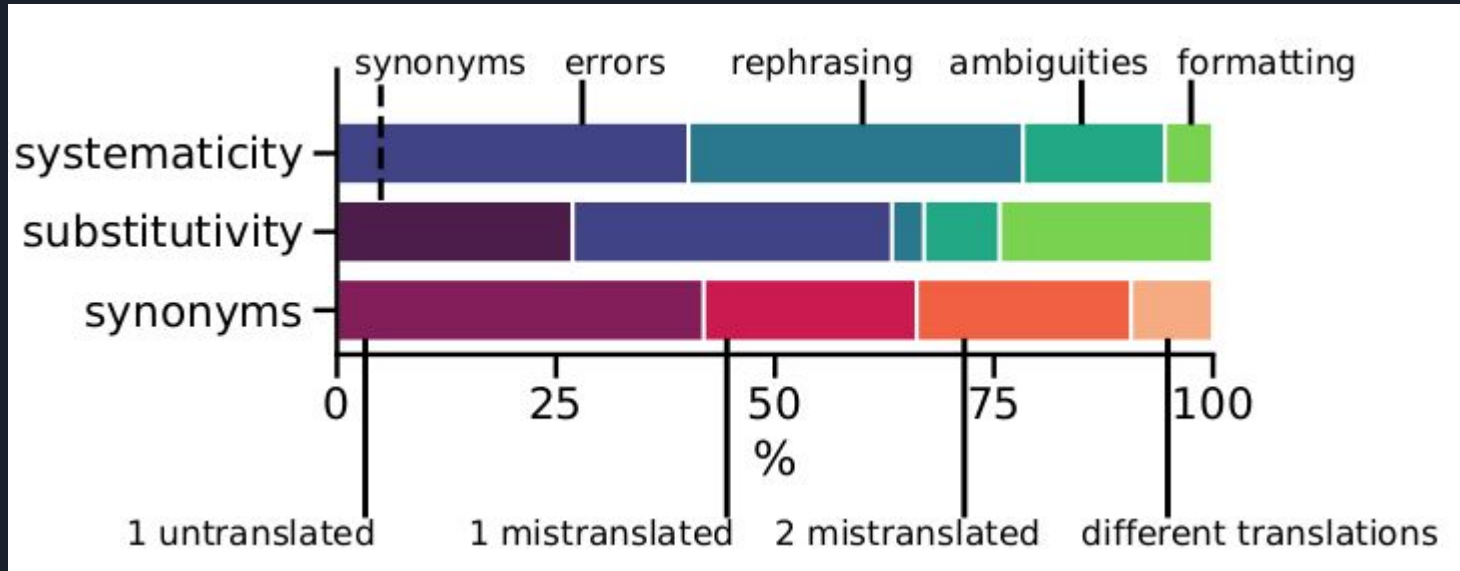
# Substitutivity

Okay, this is not as sensational as a radio-control **aeroplane**, but it's a lot quieter.

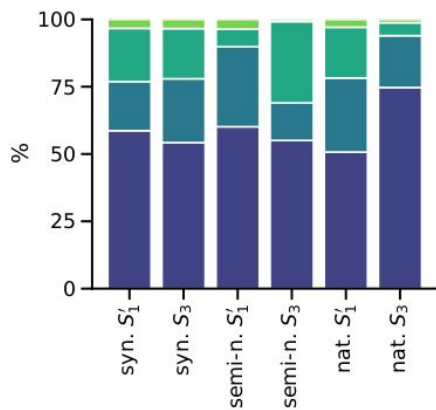
Okay, this is not as sensational as a radio-control **airplane**, but it's a lot quieter.



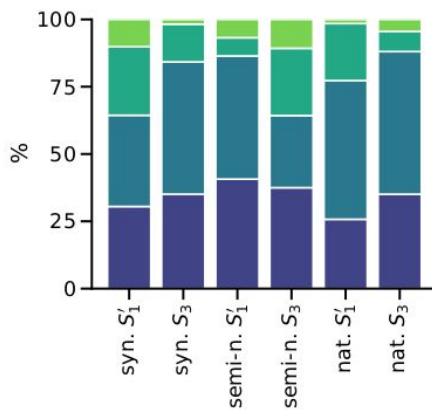
# What are those inconsistencies?



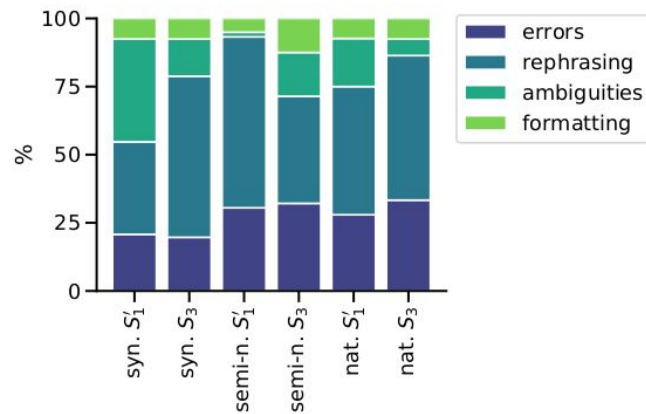
# How many of these changes are wrong?



(a) Small training set size

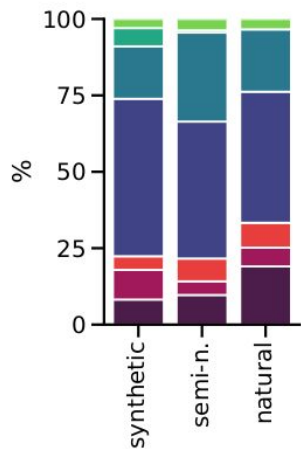


(b) Medium training set size

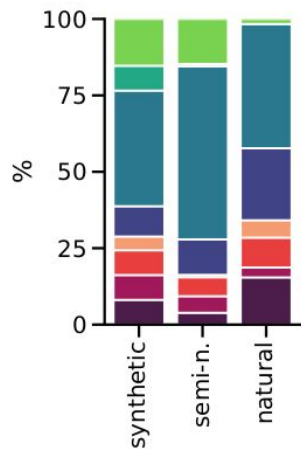


(c) Full training set

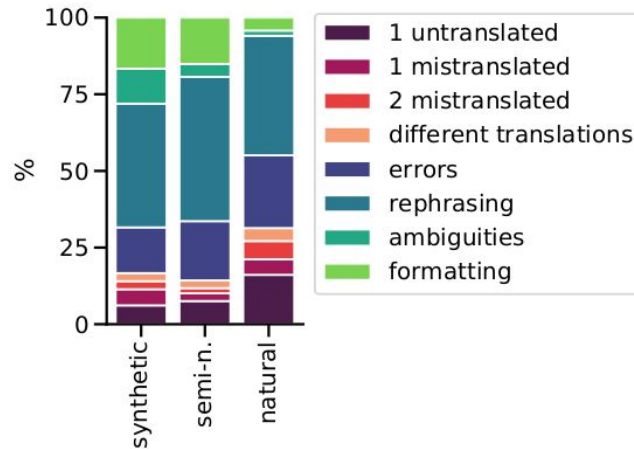
# How many of these changes are wrong?



(a) Small training set size



(b) Medium training set size



(c) Full training set



# So?

- *How do we distinguish problematic uncompositional behaviour from unproblematic uncompositional behaviour.*
- *How do we distinguish “needed” uncompositional behaviour from “unneeded” uncompositional behaviour?*
- *Are the models trained on more data really more compositional, or is compositionality just not needed as much when we have tons of data?*
- *How compositional is language, really?*



# Thank you



Verna Dankers



Elia Bruni

[dieuwkehupkes.nl](http://dieuwkehupkes.nl)  
[dieuwkehupkes@fb.com](mailto:dieuwkehupkes@fb.com)