Linguistic Intersections of Language and Gender

He, she, they, they
A first discriminative analysis of third-person pronoun semantics

Dominic Schmitz

Heinrich-Heine-Universität Düsseldorf

Motivation

Motivation

- In recent years, the use of appropriate third-person pronouns has gained increased attention

Motivation

- In recent years, the use of appropriate third-person pronouns has gained increased attention
- In English linguistics, this attention overwhelmingly manifests in form of sociolinguistic and syntactic research (e.g. Conrod, 2020; Han & Moulton, 2022; Konnelly et al., 2020)

Motivation

- In recent years, the use of appropriate third-person pronouns has gained increased attention
- In English linguistics, this attention overwhelmingly manifests in form of sociolinguistic and syntactic research (e.g. Conrod, 2020; Han & Moulton, 2022; Konnelly et al., 2020)
- What is missing, however, is
  - a semantic account of not only third-person pronouns,
  - but pronouns in general

Motivation

- In recent years, the use of appropriate third-person pronouns has gained increased attention
- In English linguistics, this attention overwhelmingly manifests in form of sociolinguistic and syntactic research (e.g. Conrod, 2020; Han & Moulton, 2022; Konnelly et al., 2020)
- What is missing, however, is
  - a semantic account of not only third-person pronouns,
  - but pronouns in general
- The present pilot study offers a first account of pronoun semantics by example of *he*, *she*, and plural and singular *they*

Singular *they*

Singular *they*

- In contemporary English, one can differentiate at least four types of singular *they* (Conrod, 2020)

Singular *they*

- In contemporary English, one can differentiate at least four types of singular *they* (Conrod, 2020)
  - **generic indefinite**

    Someone ran out of the classroom, but **they** forgot **their** backpack.

Singular *they*

- In contemporary English, one can differentiate at least four types of singular *they* (Conrod, 2020)
  - **generic indefinite**

    Someone ran out of the classroom, but **they** forgot **their** backpack.
  - **generic definite**

    The **ideal student** completes the homework, but not if **they** have an emergency.

Singular *they*

- In contemporary English, one can differentiate at least four types of singular *they* (Conrod, 2020)
  - **generic indefinite**

    Someone ran out of the classroom, but **they** forgot **their** backpack.
  - **generic definite**

    The **ideal student** completes the homework, but not if **they** have an emergency.
  - **specific definite ungendered**

    The **math teacher** is talented, but **they** hand back grades late.

Singular *they*

- In contemporary English, one can differentiate at least four types of singular *they* (Conrod, 2020)
  - **generic indefinite**

    Someone ran out of the classroom, but **they** forgot **their** backpack.
  - **generic definite**

    The **ideal student** completes the homework, but not if **they** have an emergency.
  - **specific definite ungendered**

    The **math teacher** is talented, but **they** hand back grades late.
  - **specific definite gendered**

    James is great at laundry, but **they** never wash their dishes.

Singular *they*

- In contemporary English, one can differentiate at least four types of singular *they* (Conrod, 2020)
  - **generic indefinite**

    Someone ran out of the classroom, but **they** forgot **their** backpack.
  - **generic definite**

    The **ideal student** completes the homework, but not if **they** have an emergency.
  - **specific definite ungendered**

    The **math teacher** is talented, but **they** hand back grades late.
  - **specific definite gendered**

    James is great at laundry, but **they** never wash their dishes.

Research questions

Research questions

RQ1 – Methodological Question

Can distributional semantics meaningfully capture pronoun semantics?

Research questions

RQ1 – Methodological Question

Can distributional semantics meaningfully capture pronoun semantics?

RQ2 – Theoretical Question

How is singular *they* semantically related to other third-person pronouns?

Methods

discriminative learning and instance vectors

General idea

General idea

- Simulate an individual's mental lexicon by implementing a linear discriminative learning network (e.g. Baayen et al., 2019)

General idea

- Simulate an individual's mental lexicon by implementing a linear discriminative learning network (e.g. Baayen et al., 2019)
- For this, semantics and word forms are required as starting points

General idea

- Simulate an individual's mental lexicon by implementing a linear discriminative learning network (e.g. Baayen et al., 2019)
- For this, semantics and word forms are required as starting points

corpus

General idea

- Simulate an individual's mental lexicon by implementing a linear discriminative learning network (e.g. Baayen et al., 2019)
- For this, semantics and word forms are required as starting points

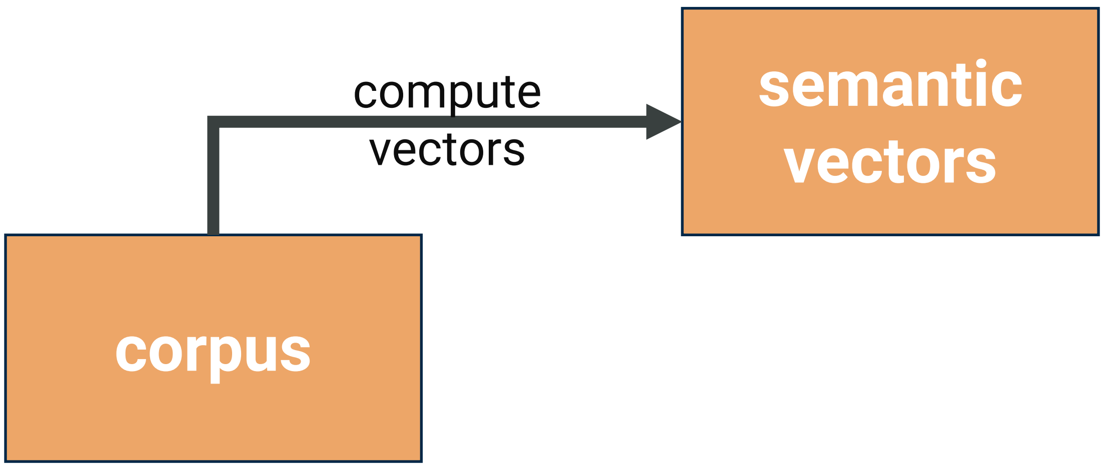

General idea

- Simulate an individual's mental lexicon by implementing a linear discriminative learning network (e.g. Baayen et al., 2019)
- For this, semantics and word forms are required as starting points

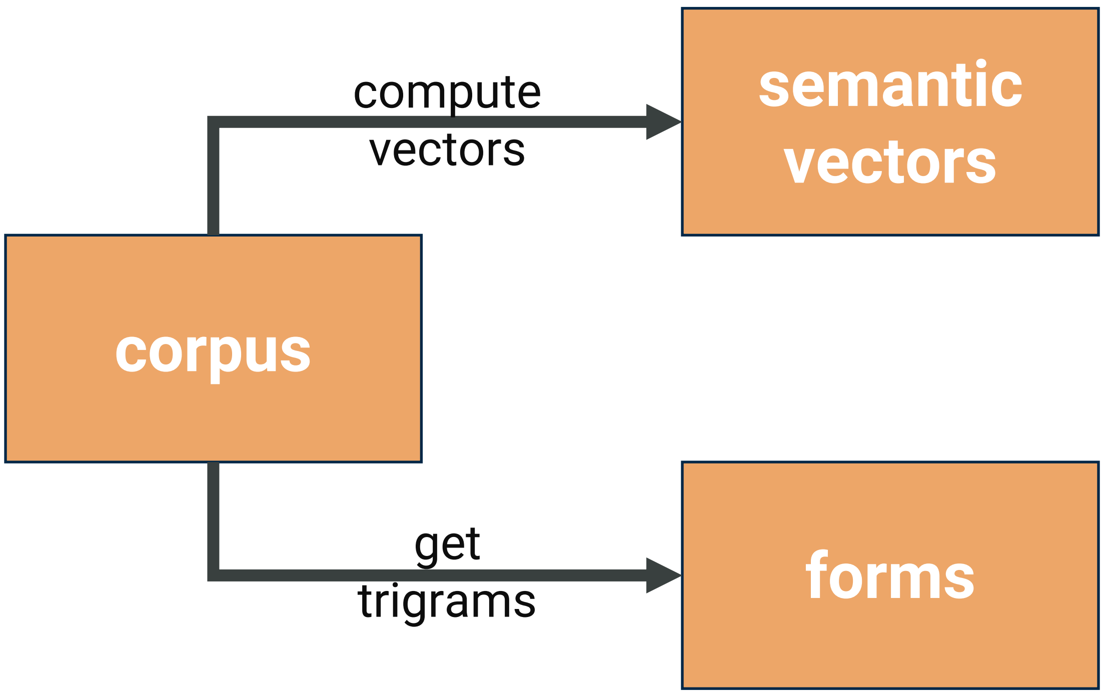

General idea

- Simulate an individual's mental lexicon by implementing a linear discriminative learning network (e.g. Baayen et al., 2019)
- For this, semantics and word forms are required as starting points

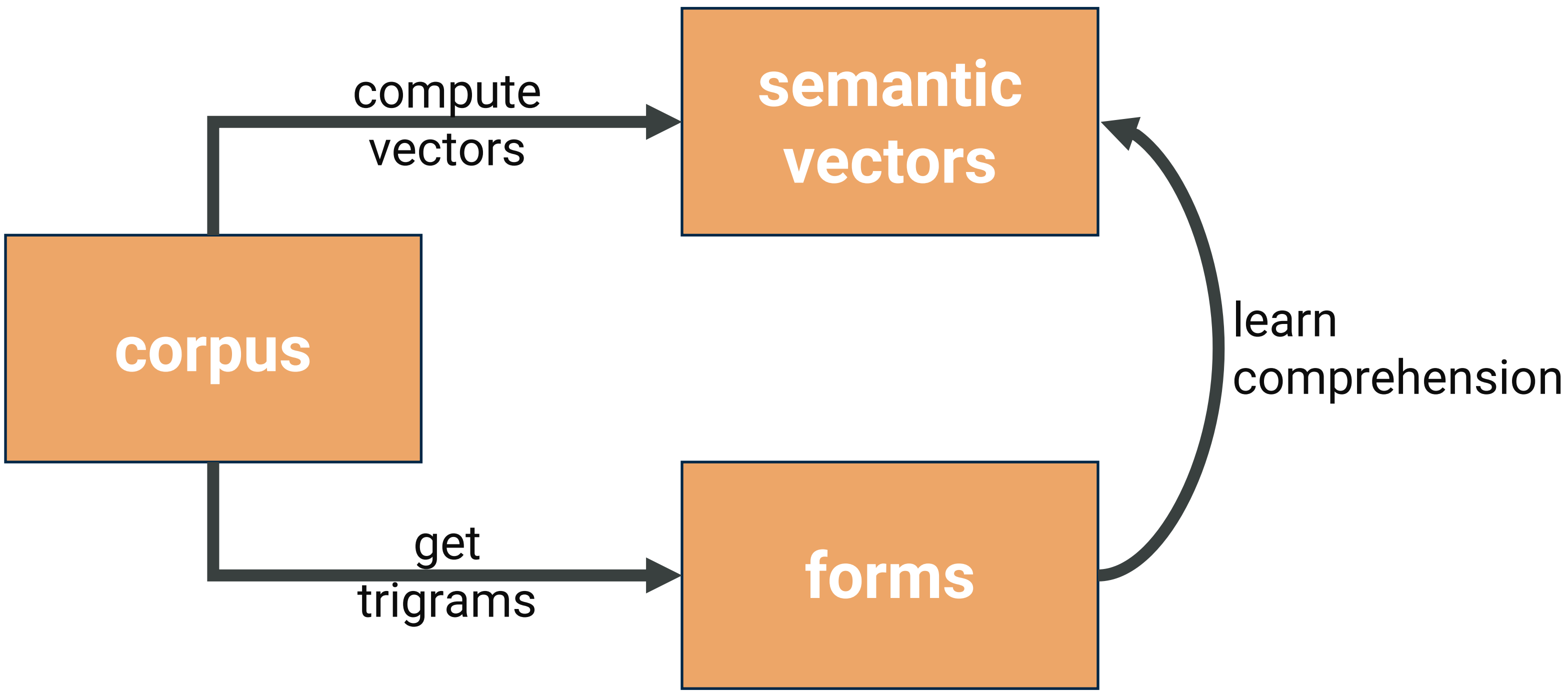

Corpus

Corpus

- Small corpus based on COCA (Davies, 2008-)
  - 17,805 word form tokens
  - 1,000 sentences
  - 30 + attestations of each target pronoun

    he, *she*, and plural and singular *they*

Corpus

- Small corpus based on COCA (Davies, 2008-)
  - 17,805 word form tokens
  - 1,000 sentences
  - 30 + attestations of each target pronoun

    he, *she*, and plural and singular *they*
- Pronoun attestations were manually checked for number and genericity

Corpus

- Small corpus based on COCA (Davies, 2008-)
  - 17,805 word form tokens
  - 1,000 sentences
  - 30 + attestations of each target pronoun
    he, *she*, and plural and singular *they*
- Pronoun attestations were manually checked for number and genericity
- Automatically analysed and annotated for inflection using the RNNTagger software (Schmid, 1999)

General idea

- Simulate an individual's mental lexicon by implementing a linear discriminative learning network (e.g. Baayen et al., 2019)
- For this, semantics and word forms are required as starting points

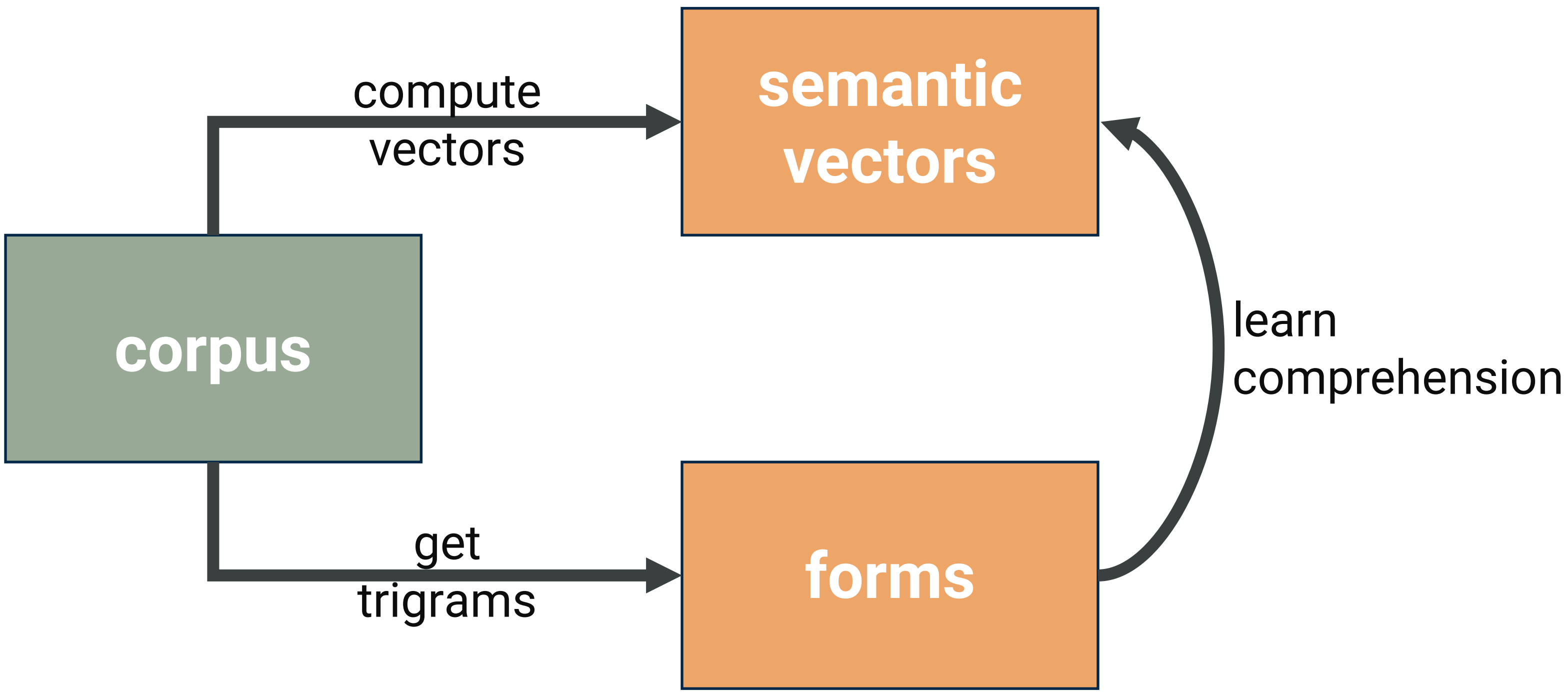

Semantic vectors

Semantic vectors

- Distributional Hypothesis (e.g. Harris, 1954)

difference in meaning ↔ difference in distribution

Semantic vectors

- Distributional Hypothesis (e.g. Harris, 1954)

  difference in meaning ↔ difference in distribution

- Difference in meaning is measured via semantic vectors

Semantic vectors

- Distributional Hypothesis (e.g. Harris, 1954)

  difference in meaning ↔ difference in distribution

- Difference in meaning is measured via semantic vectors
- There are different algorithms to arrive at a word's semantic vector, two of them are
  - NDL: Naive Discriminative Learning (Baayen et al., 2011)
  - Instance vectors (Lapesa et al., 2018)

Naive Discriminative Learning

Naive Discriminative Learning

- Taking the whole corpus, semantic vectors for bases and inflectional functions were computed using NDL

Naive Discriminative Learning

- Taking the whole corpus, semantic vectors for bases and inflectional functions were computed using NDL
- NDL follows the Rescorla-Wagner rules (Rescorla & Wagner, 1972; Wagner & Rescorla, 1972)

Naive Discriminative Learning

- Taking the whole corpus, semantic vectors for bases and inflectional functions were computed using NDL
- NDL follows the Rescorla-Wagner rules (Rescorla & Wagner, 1972; Wagner & Rescorla, 1972)
- Most importantly, these rules state that
  - Outcomes are predicted by cues
  - The associative strength between an outcome and a cue is represented by a single number

Naive Discriminative Learning

- Taking the whole corpus, semantic vectors for bases and inflectional functions were computed using NDL
- NDL follows the Rescorla-Wagner rules (Rescorla & Wagner, 1972; Wagner & Rescorla, 1972)
- Most importantly, these rules state that
  - Outcomes are predicted by cues
  - The associative strength between an outcome and a cue is represented by a single number
- Each sentence was used to predict each individual outcome within the sentence by the other bases/function words/inflectional functions in that sentence

Naive Discriminative Learning

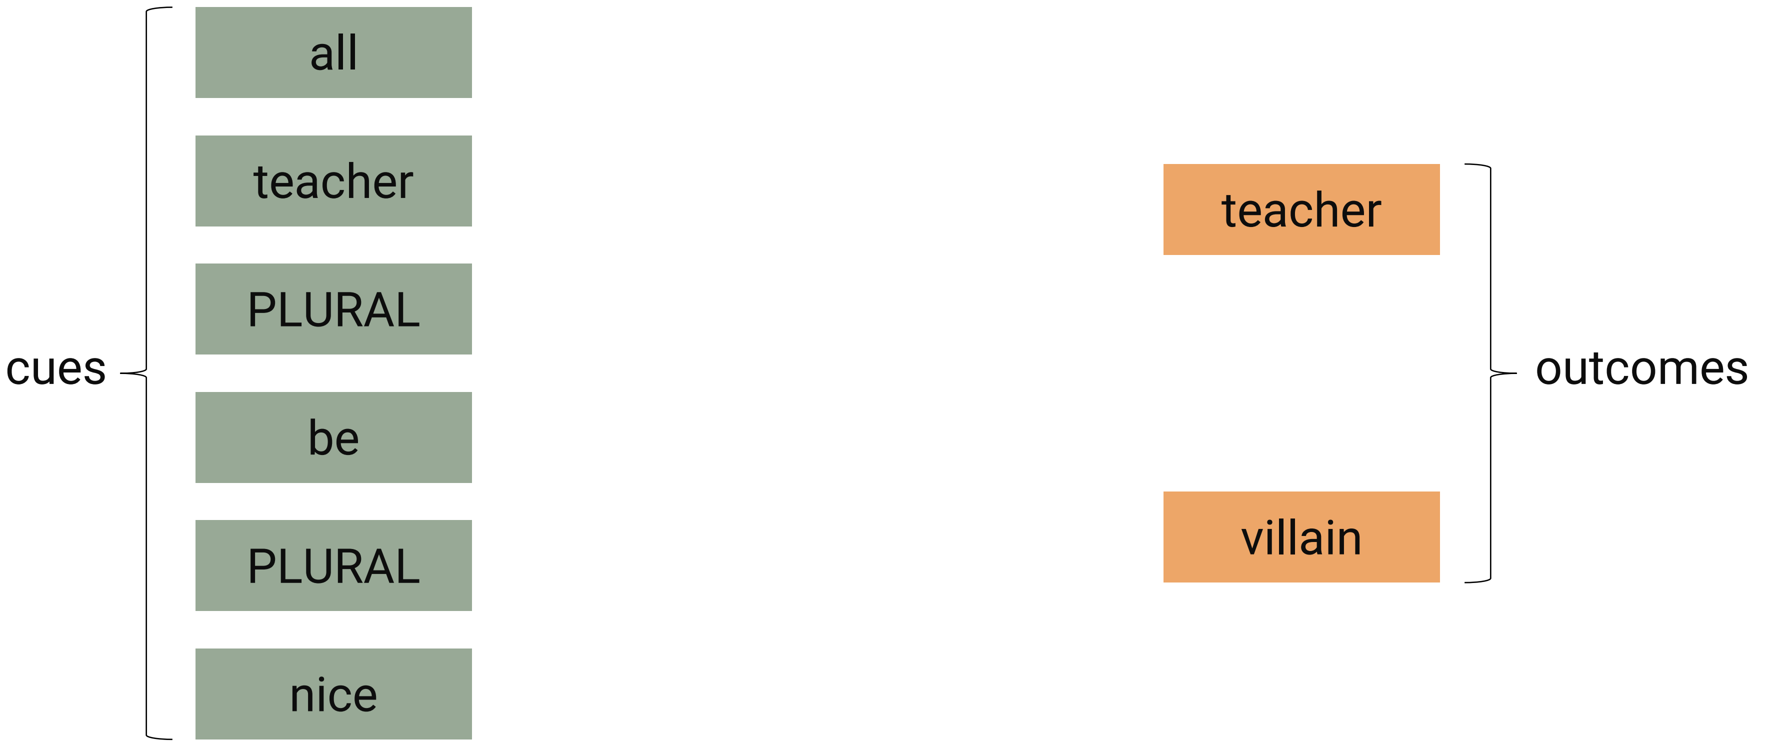

All teachers are nice.

	all	teacher	PLURAL	be	nice	villain	evil
teacher							
villain							

Naive Discriminative Learning

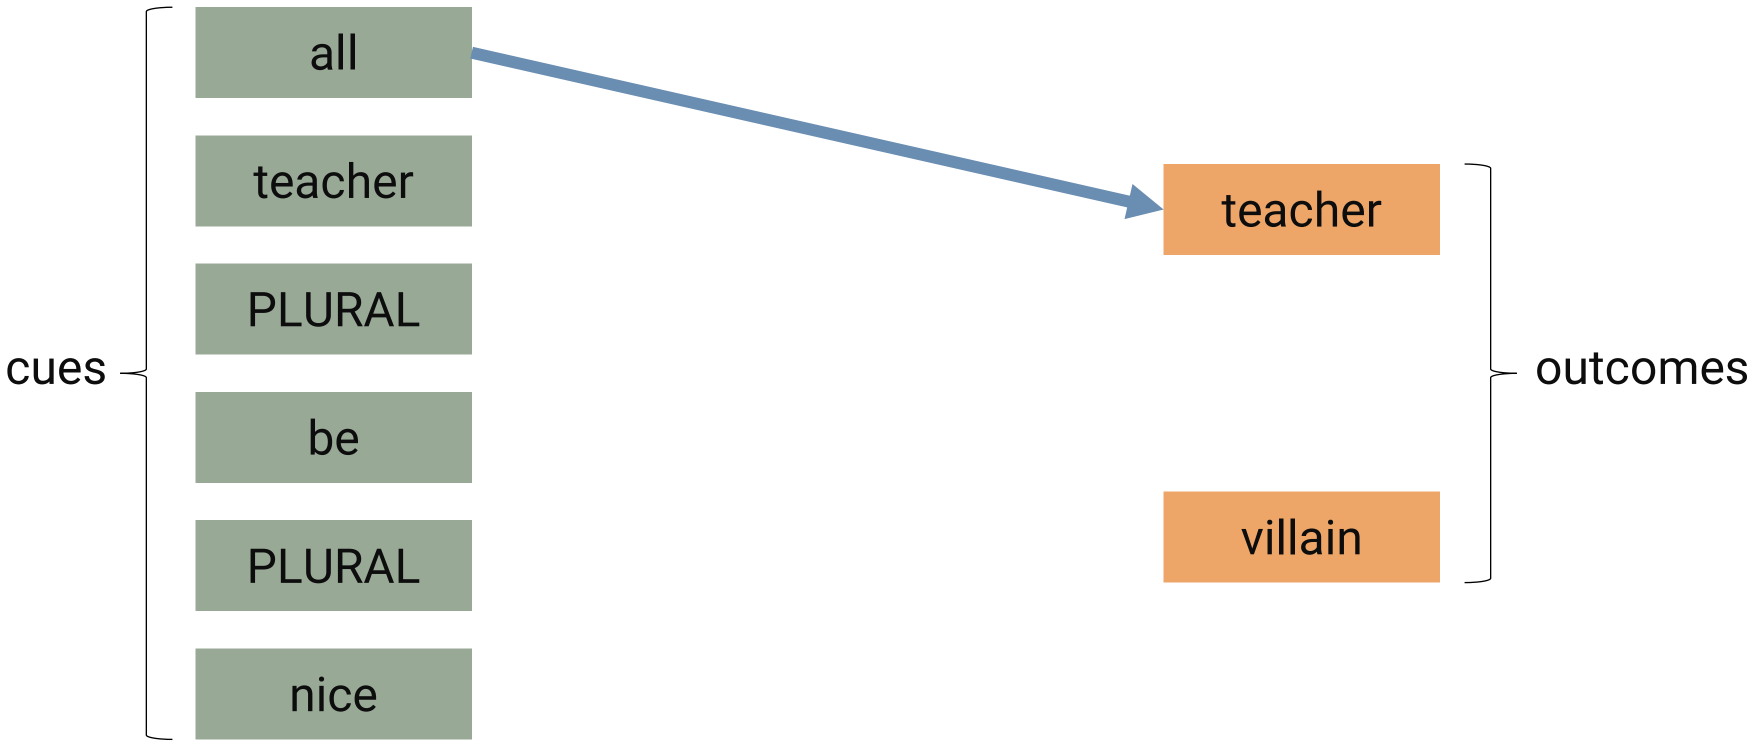

All teachers are nice.

	all	teacher	PLURAL	be	nice	villain	evil
teacher	+						
villain							

Naive Discriminative Learning

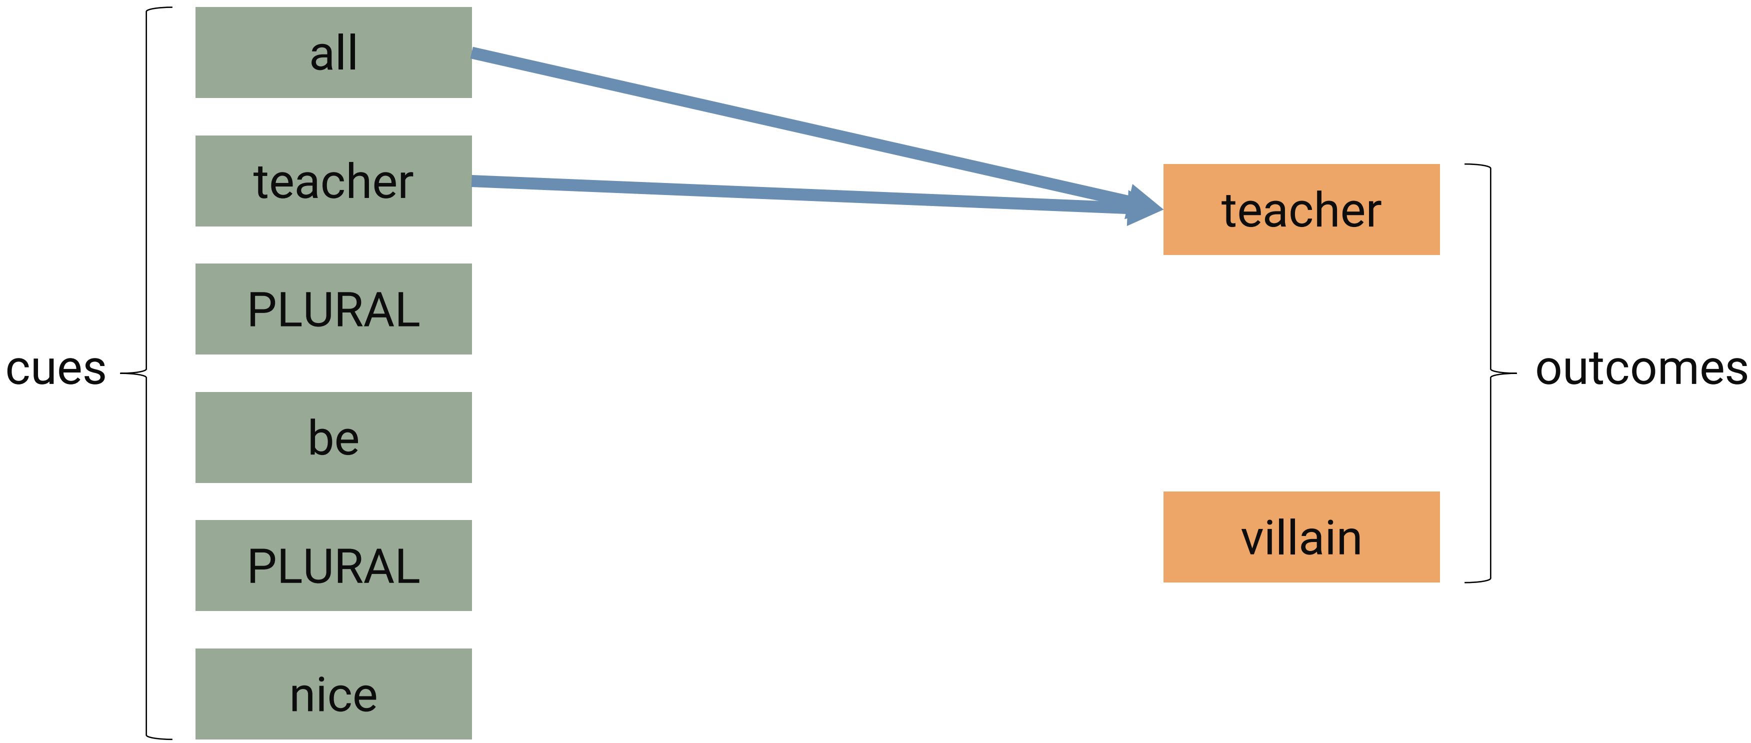

All teachers are nice.

	all	teacher	PLURAL	be	nice	villain	evil
teacher	+	+					
villain							

Naive Discriminative Learning

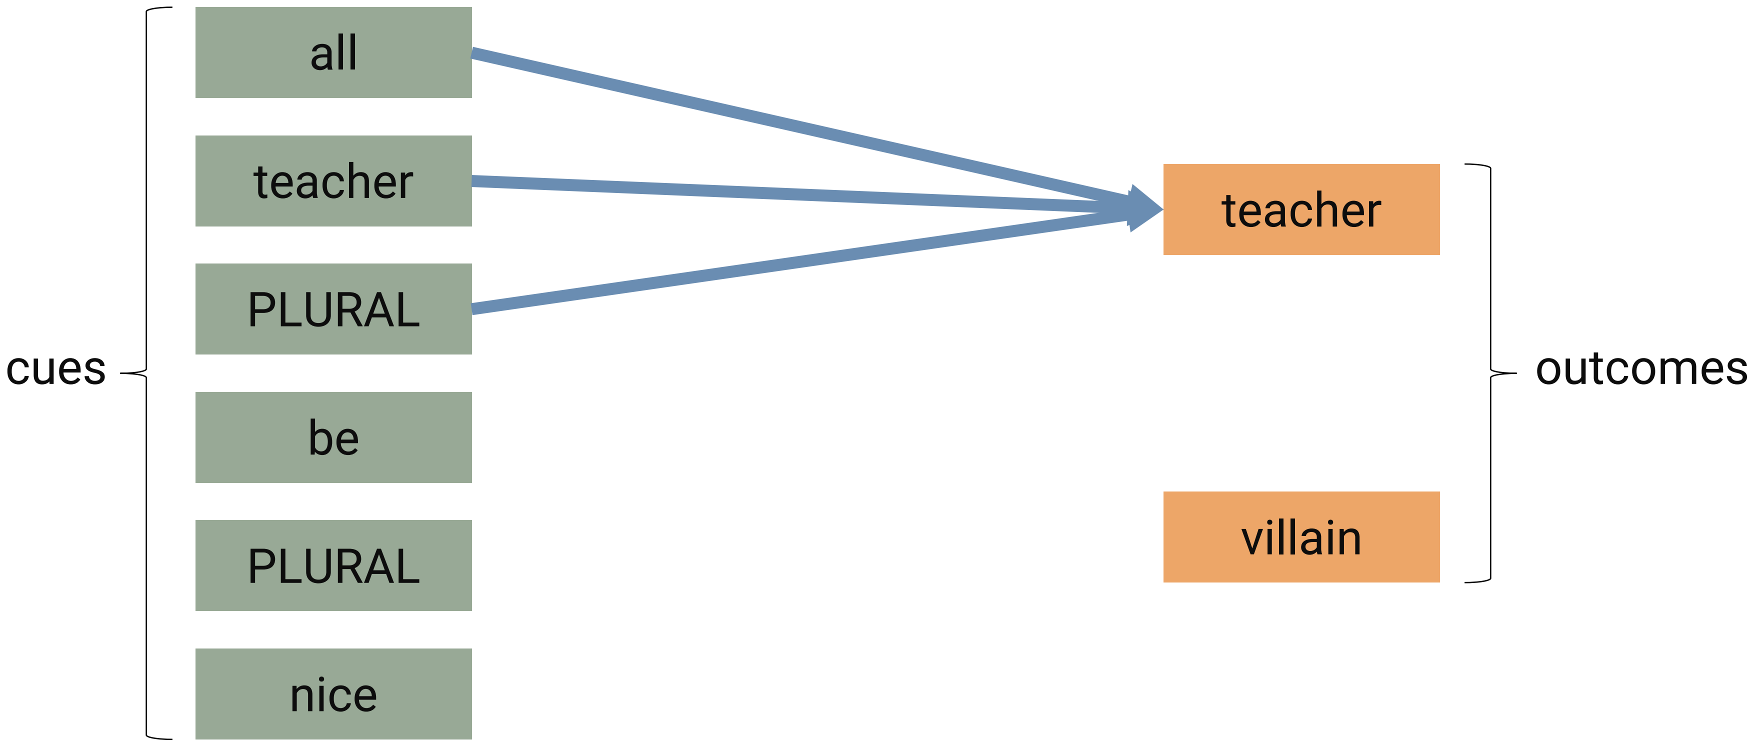

All teachers are nice.

	all	teacher	PLURAL	be	nice	villain	evil
teacher	+	+	+				
villain							

Naive Discriminative Learning

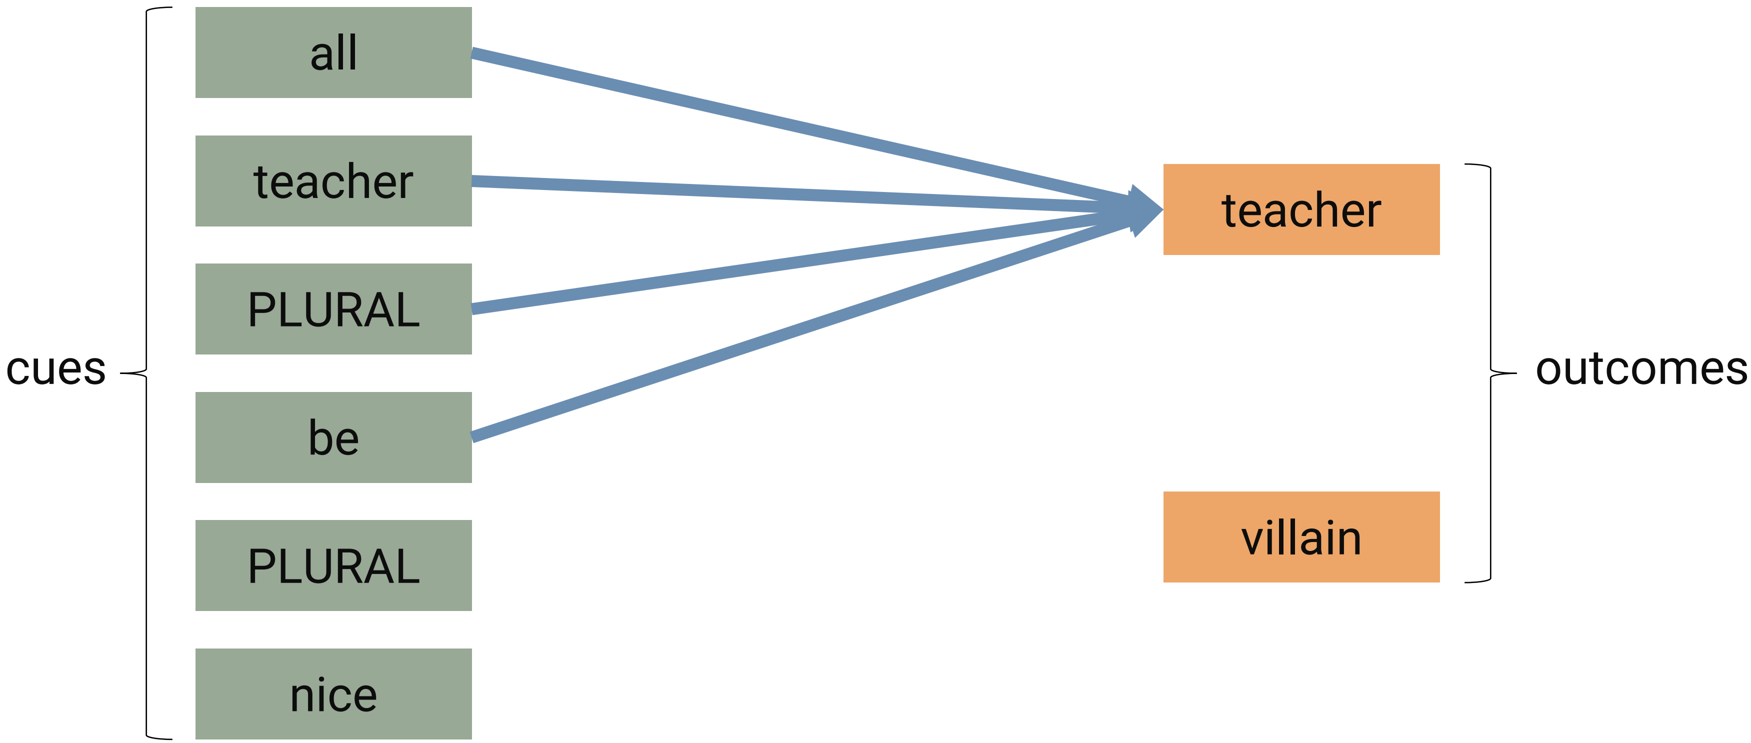

All teachers are nice.

	all	teacher	PLURAL	be	nice	villain	evil
teacher	+	+	+	+			
villain							

Naive Discriminative Learning

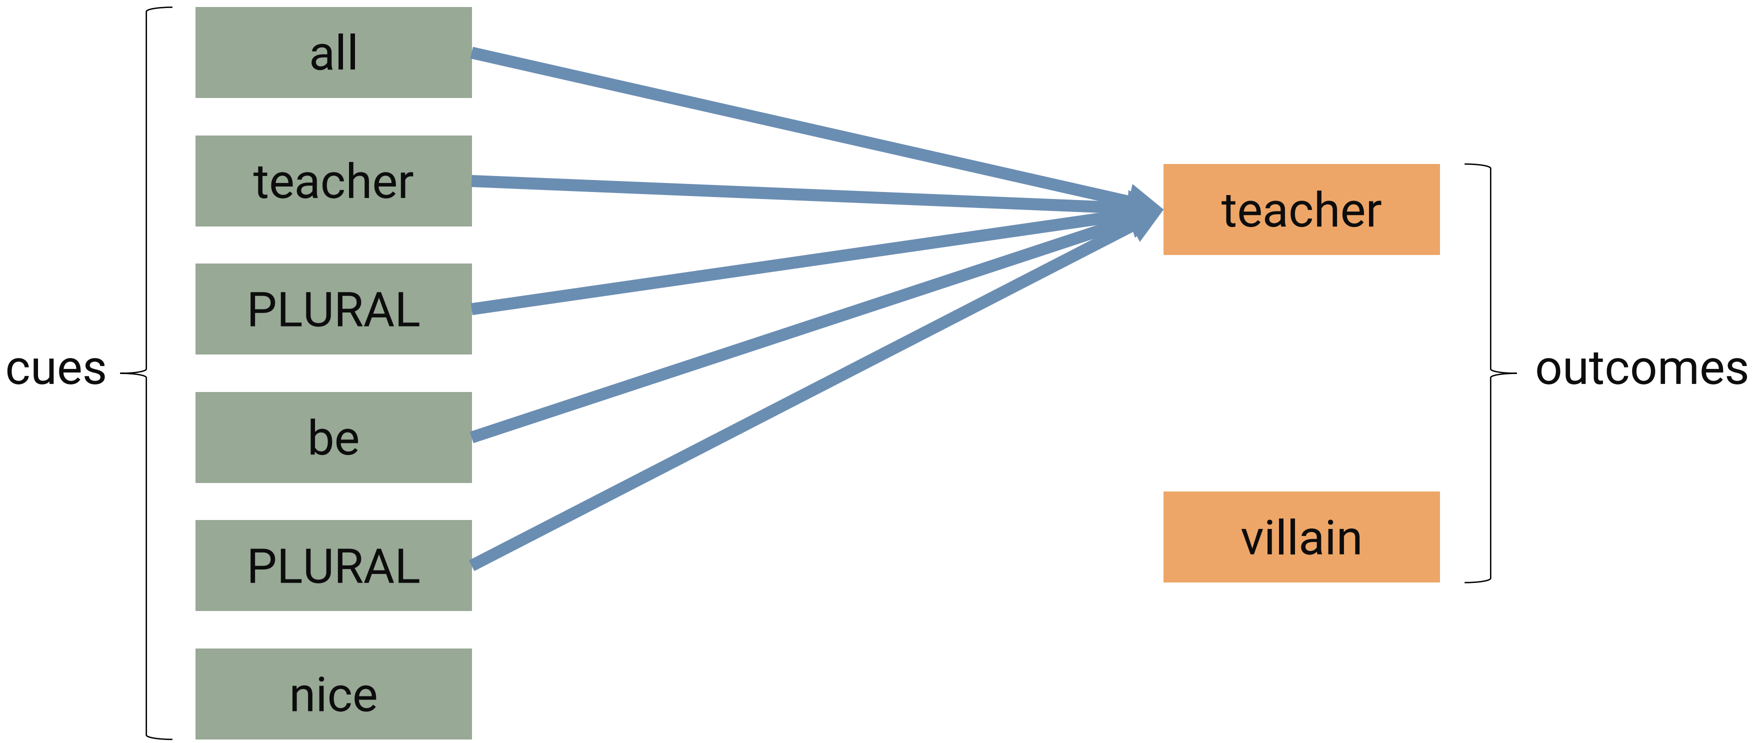

All teachers are nice.

	all	teacher	PLURAL	be	nice	villain	evil
teacher	+	+	++	+			
villain							

Naive Discriminative Learning

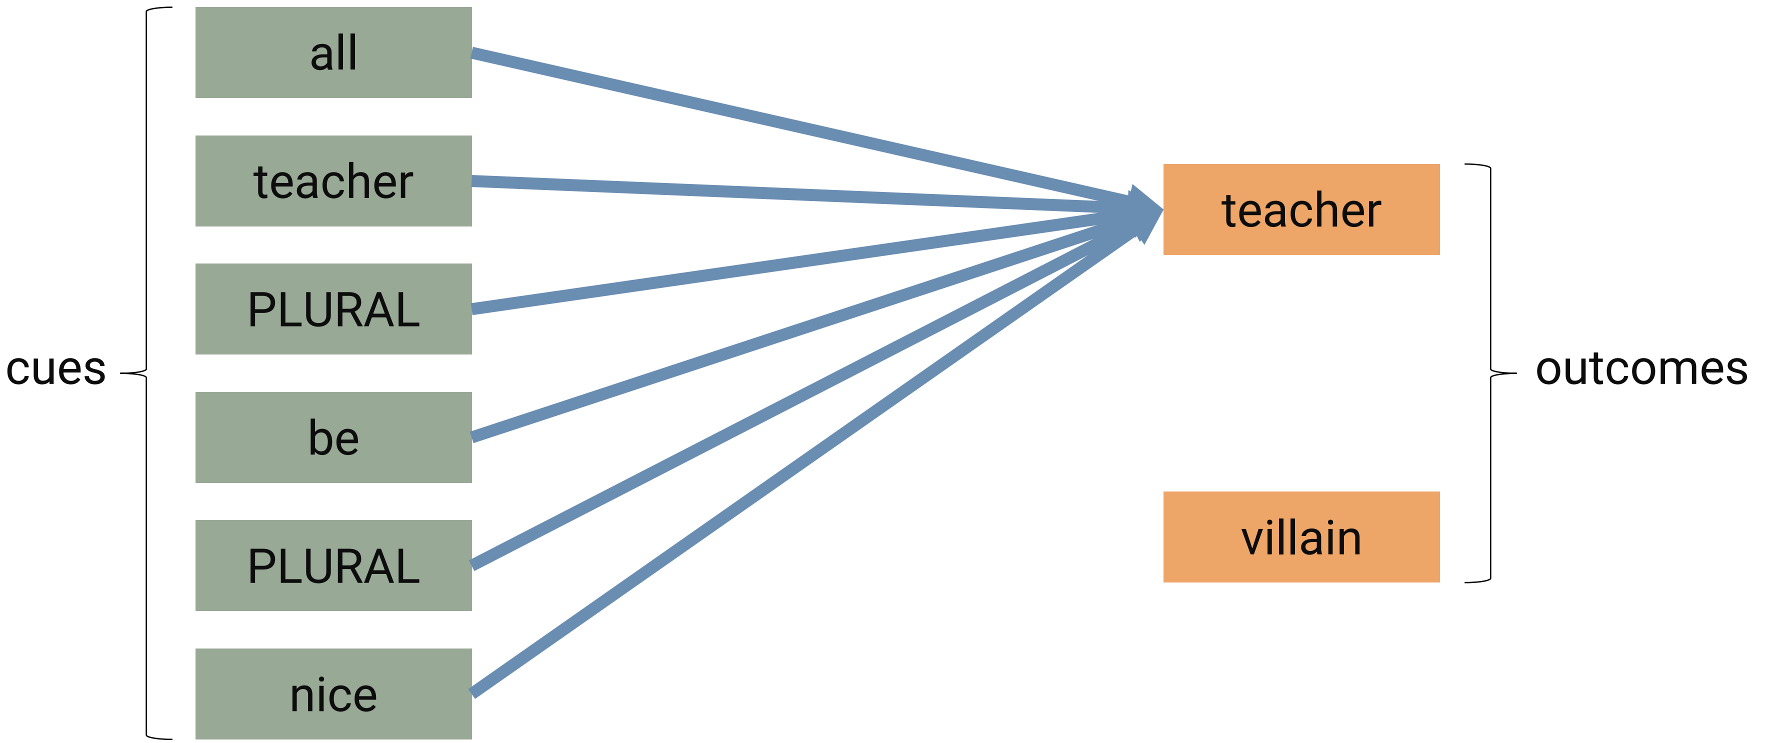

All teachers are nice.

	all	teacher	PLURAL	be	nice	villain	evil
teacher	+	+	++	+	+		
villain							

Naive Discriminative Learning

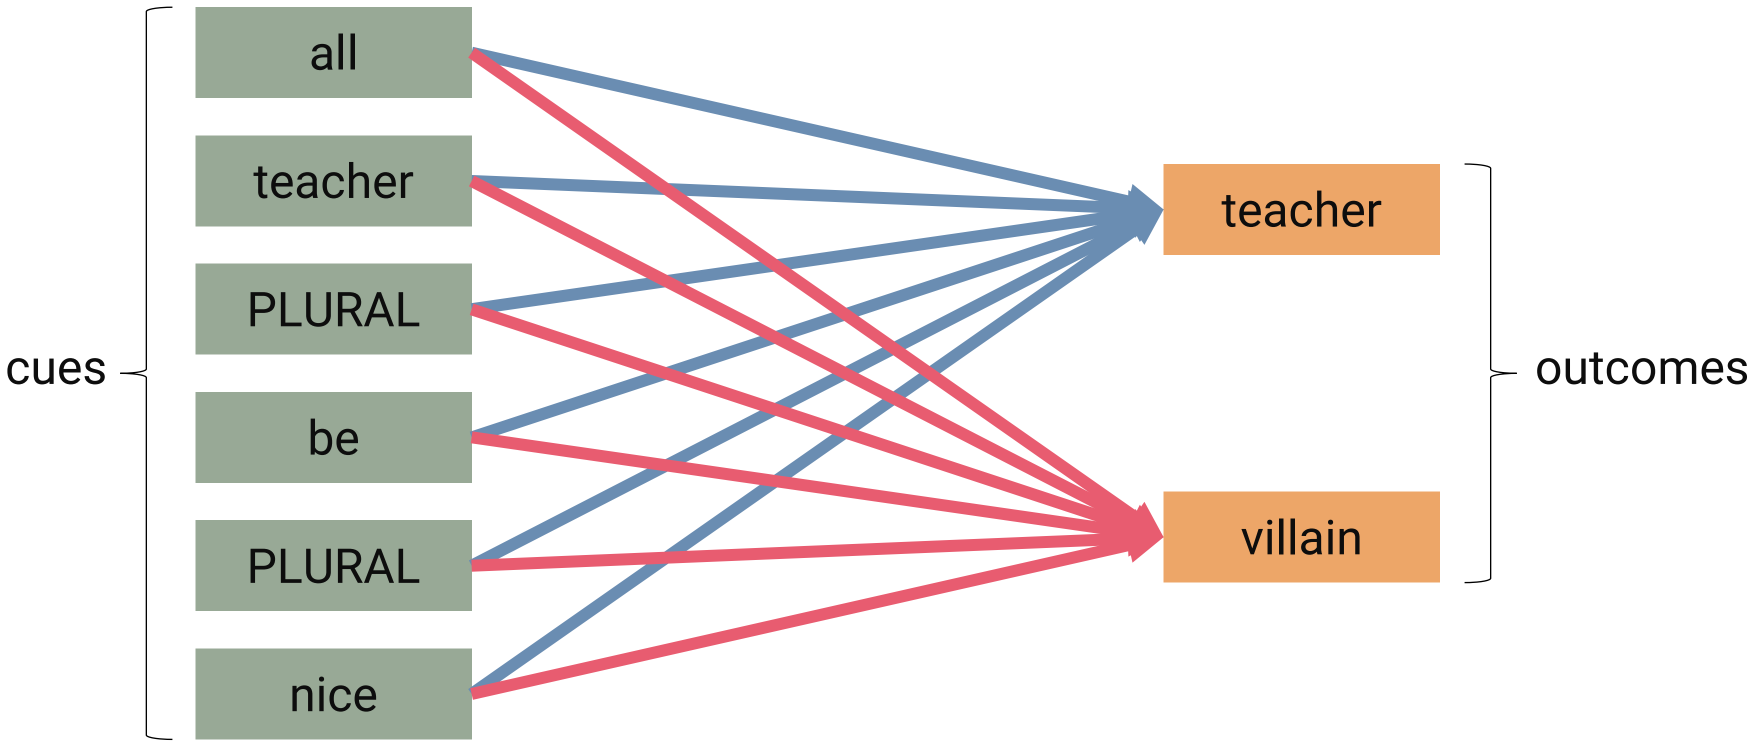

All teachers are nice.

	all	teacher	PLURAL	be	nice	villain	evil
teacher	+	+	++	+	+	-	-
villain	-	-	-	-	-		

Naive Discriminative Learning

Naive Discriminative Learning

- Result: association weights for all word bases, function words, inflectional functions

Naive Discriminative Learning

- Result: association weights for all word bases, function words, inflectional functions
- Each row of association weights corresponds to the semantic vectors of a base, function word, or inflection function

Naive Discriminative Learning

- Result: association weights for all word bases, function words, inflectional functions
- Each row of association weights corresponds to the semantic vectors of a base, function word, or inflection function
- For example:

	apple	drink	fork	potato	universe	star
eat	0.3	0.2	0.5	0.4	0.00002	0.000071
astronomy	0.0003	0.0015	0.00704	0.0003	0.6	0.8

Naive Discriminative Learning

- Result: association weights for all word bases, function words, inflectional functions
- Each row of association weights corresponds to the semantic vectors of a base, function word, or inflection function
- For example:

	apple	drink	fork	potato	universe	star
eat	0.3	0.2	0.5	0.4	0.00002	0.000071
astronomy	0.0003	0.0015	0.00704	0.0003	0.6	0.8

Naive Discriminative Learning

Naive Discriminative Learning

- However: 1 vector per base/function word/inflectional function

  = 1 vector per pronoun

Naive Discriminative Learning

- However: 1 vector per base/function word/inflectional function
  = 1 vector per pronoun
- Potentially very different semantics of pronoun attestations are conflated into one vector representation

Naive Discriminative Learning

- However: 1 vector per base/function word/inflectional function

  = 1 vector per pronoun

- Potentially very different semantics of pronoun attestations are conflated into one vector representation

- This is an issue!

  → Pronouns are assumed to inherit the semantics of their referents

Instance vectors

- preceding and following words

Instance vectors

- The solution: instance vectors
  - preceding and following words

Instance vectors

- The solution: instance vectors
  - Take n preceding and following words

Instance vectors

- The solution: instance vectors
  - Take n preceding and following words
  - Get the semantic vectors of these words

Instance vectors

- The solution: instance vectors
  - Take n preceding and following words
  - Get the semantic vectors of these words
  - Compute the mean of these vectors

    = instance vector

Instance vectors

- The solution: instance vectors
  - Take n preceding and following words
  - Get the semantic vectors of these words
  - Compute the mean of these vectors

    = instance vector

Instance vectors

- The solution: instance vectors
  - Take n preceding and following words
  - Get the semantic vectors of these words
  - Compute the mean of these vectors

    = instance vector

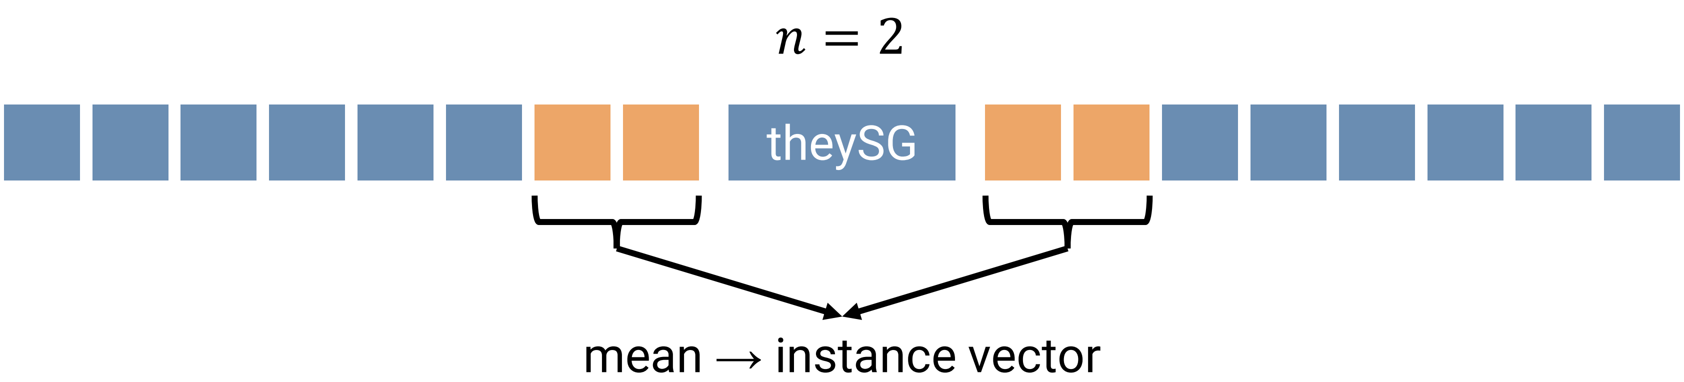

Instance vectors

- The solution: instance vectors
  - Take n preceding and following words
  - Get the semantic vectors of these words
  - Compute the mean of these vectors

    = instance vector

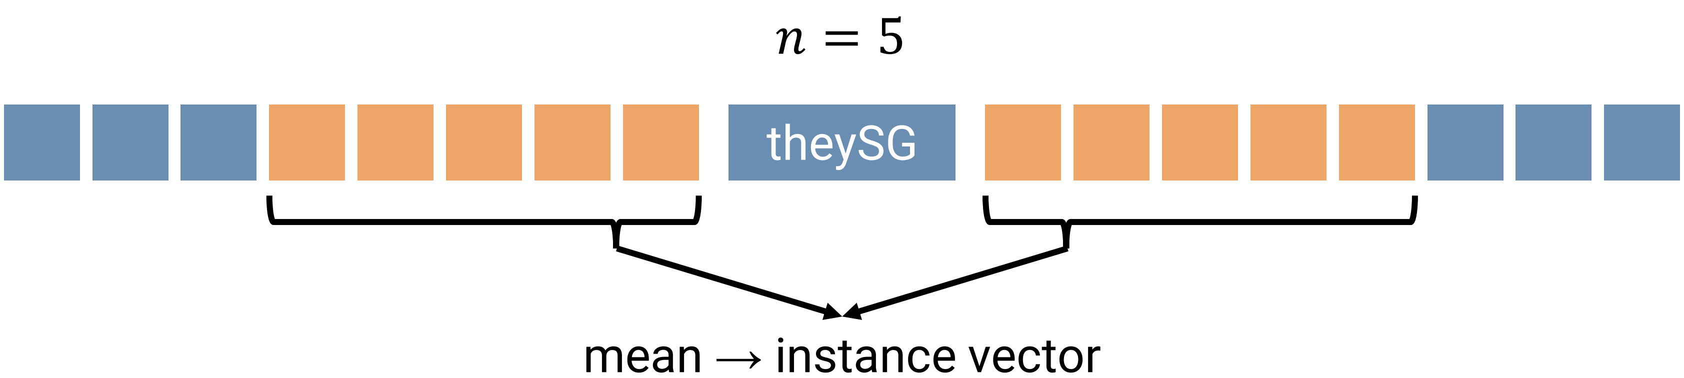

Instance vectors

- The solution: instance vectors
  - Take n preceding and following words
  - Get the semantic vectors of these words
  - Compute the mean of these vectors

    = instance vector

$n = 8$

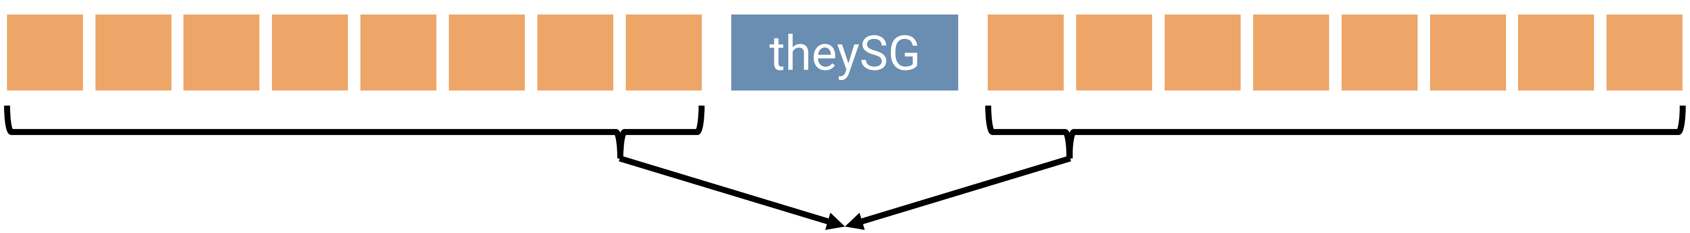

mean → instance vector

Instance vectors

- For the present study
  - $n = 5$
  - Preceding and following units: vectors for bases/function words/inflectional functions
  - Preceding and following semantic vectors: via NDL

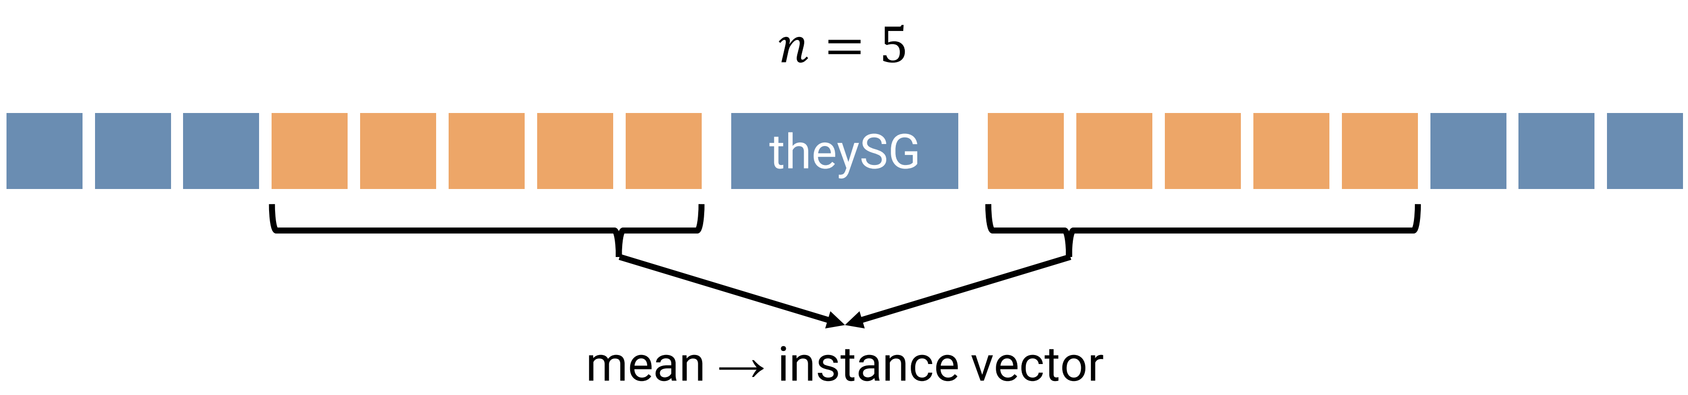

General idea

- Simulate an individual's mental lexicon by implementing a linear discriminative learning network (e.g. Baayen et al., 2019)
- For this, semantics and word forms are required as starting points

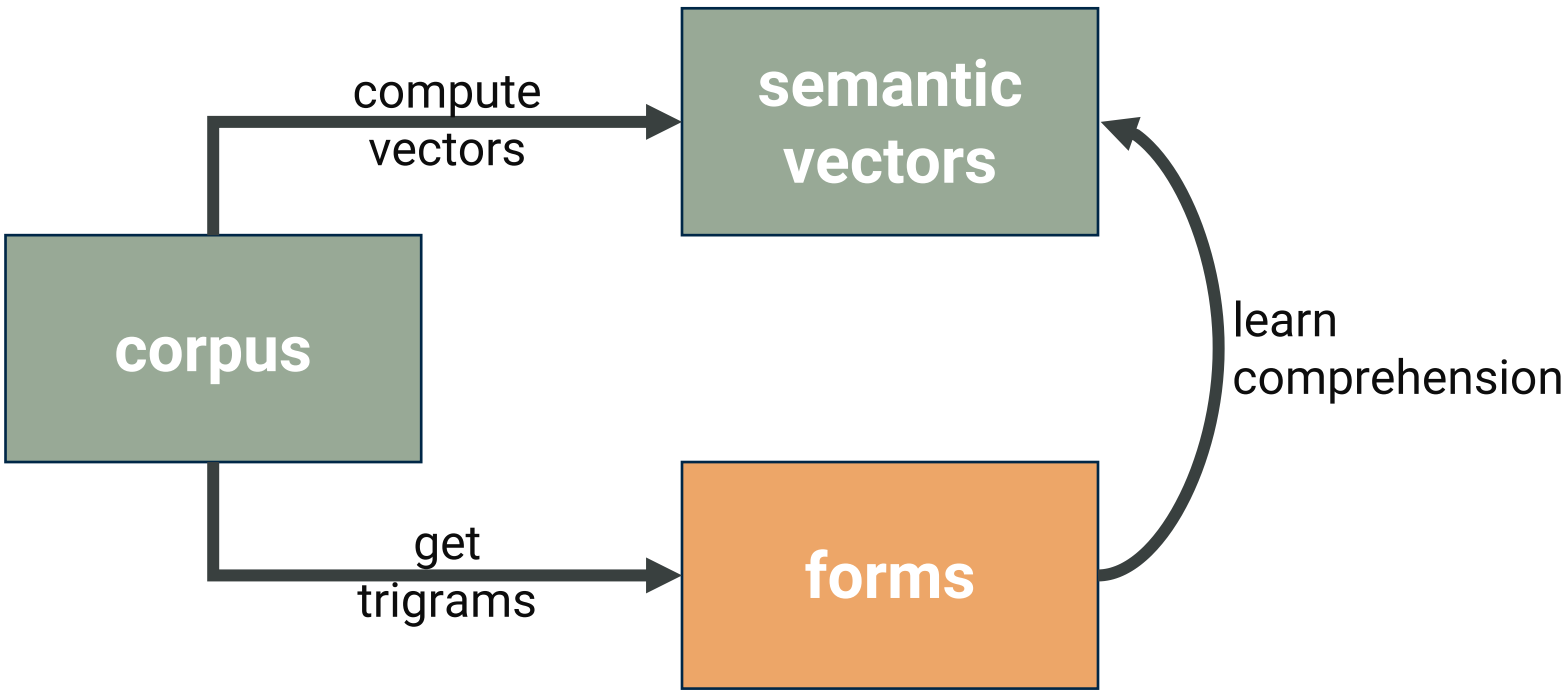

Form vectors

Form vectors

- Trigrams as unit for a word’s form

Form vectors

- Trigrams as unit for a word's form
- Trigrams / triphones have been shown to capture the form variability of words well (e.g. Chuang et al., 2020; Schmitz et al., 2021; Schmitz et al., 2023)

Form vectors

- Trigrams as unit for a word's form
- Trigrams / triphones have been shown to capture the form variability of words well (e.g. Chuang et al., 2020; Schmitz et al., 2021; Schmitz et al., 2023)

target form	#ca	cat	at#	cap	ap#	#ba	bat
cat	1	1	1	0	0	0	0
cap	1	0	0	1	1	0	0
bat	0	0	1	0	0	1	1

General idea

- Simulate an individual's mental lexicon by implementing a linear discriminative learning network (e.g. Baayen et al., 2019)
- For this, semantics and word forms are required as starting points

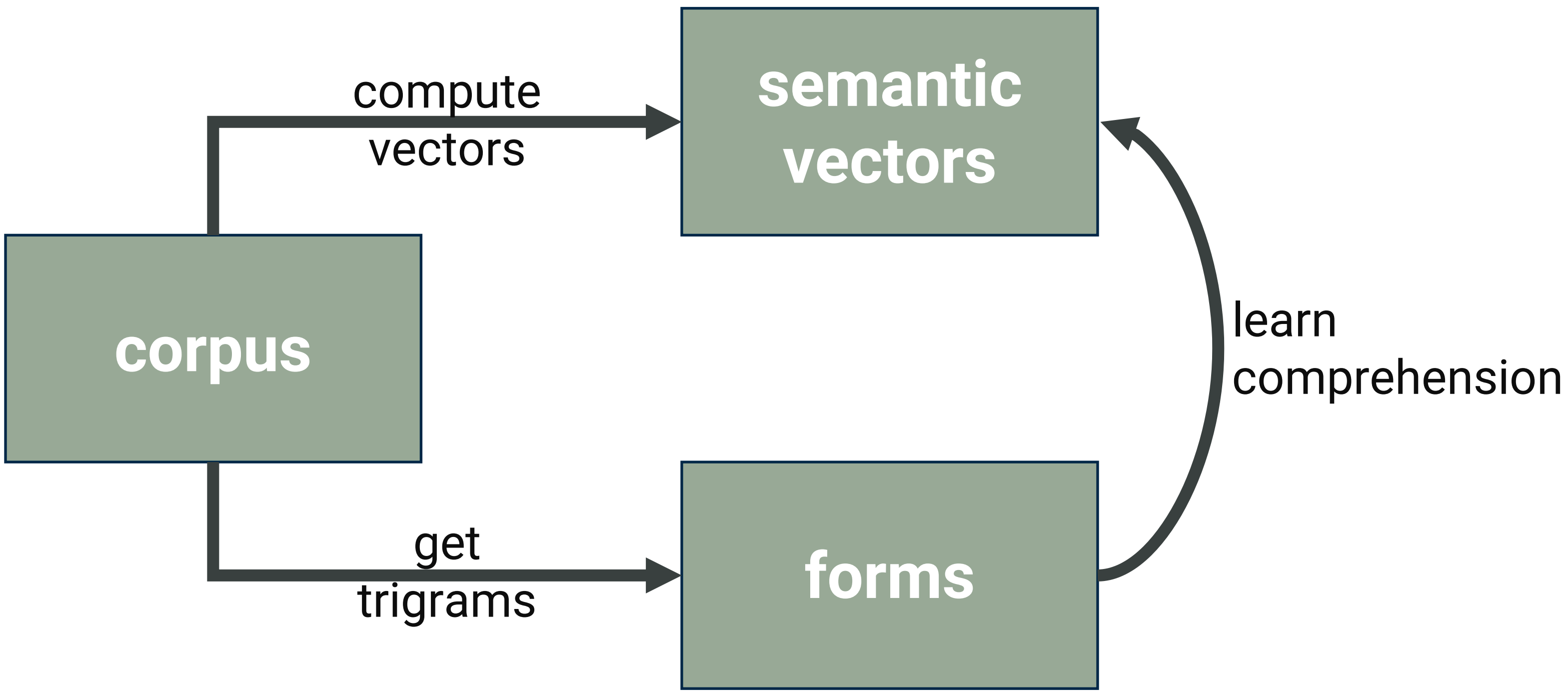

Mapping comprehension

- Comprehension is learnt by linearly mapping the matrix of forms onto the matrix of semantic vectors

Mapping comprehension

- Comprehension is learnt by linearly mapping the matrix of forms onto the matrix of semantic vectors

Mapping comprehension

- Comprehension is learnt by linearly mapping the matrix of forms onto the matrix of semantic vectors

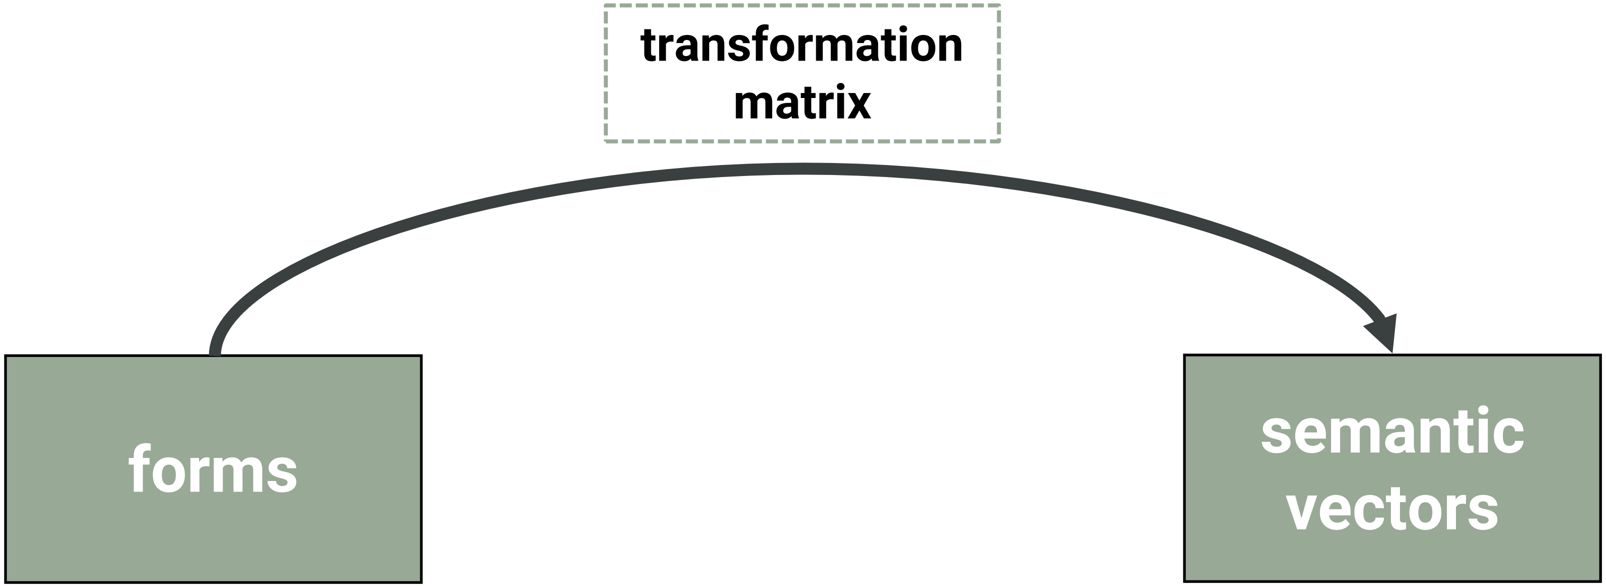

Mapping comprehension

- Comprehension is learnt by linearly mapping the matrix of forms onto the matrix of semantic vectors

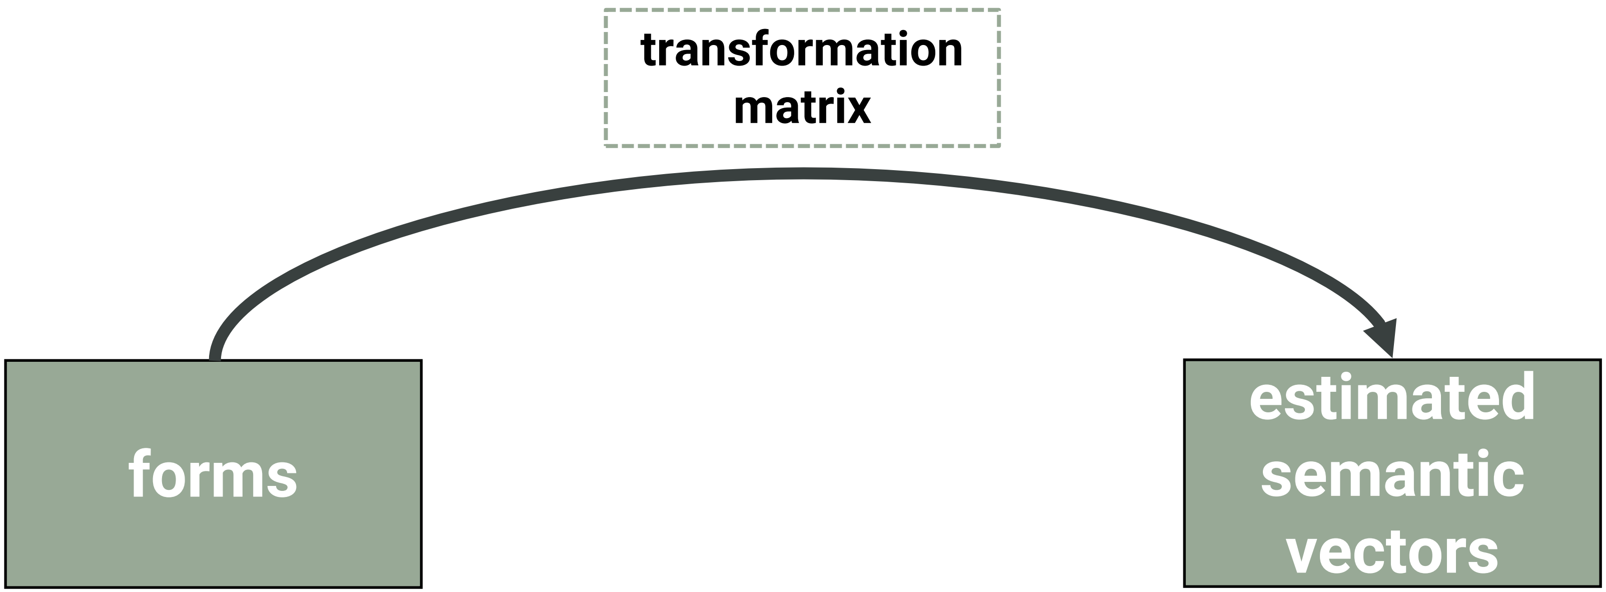

Semantic measures

- From the comprehension mapping, semantic measures can be derived

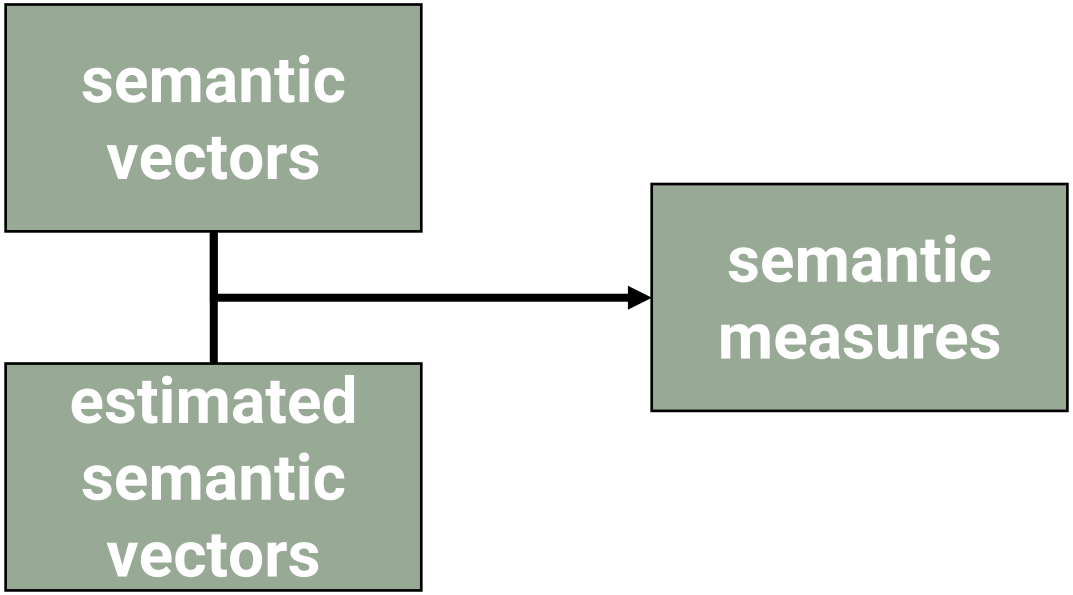

Results

Activation diversity and neighbourhood density

Semantic measures

SEMANTIC ACTIVATION DIVERSITY

Euclidian length of a target's vector

Semantic measures

SEMANTIC ACTIVATION DIVERSITY

Euclidian length of a target's vector

- Higher values indicate more co-activation in the lexicon

Semantic measures

SEMANTIC ACTIVATION DIVERSITY

Euclidian length of a target's vector

- Higher values indicate more co-activation in the lexicon

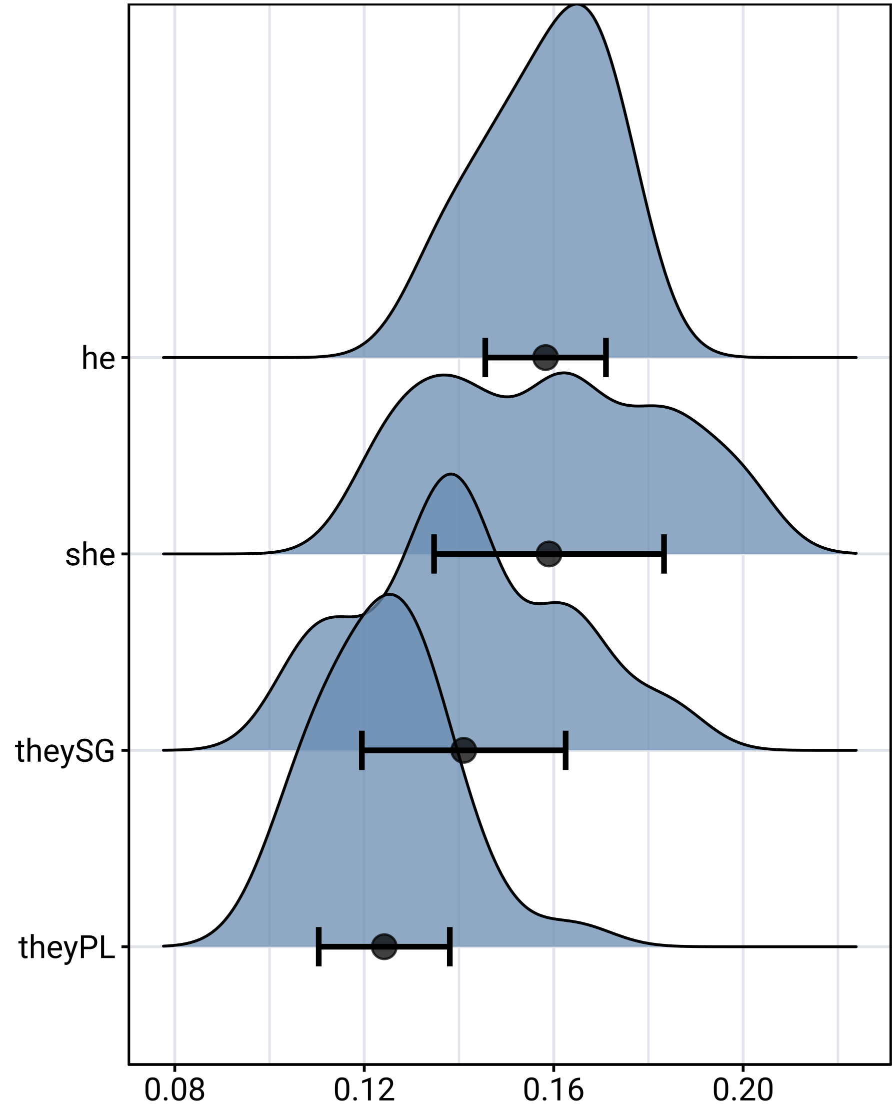

Semantic measures

SEMANTIC ACTIVATION DIVERSITY

Euclidian length of a target's vector

- Higher values indicate more co-activation in the lexicon

he, she

↓

singular *they*

↓

plural *they*

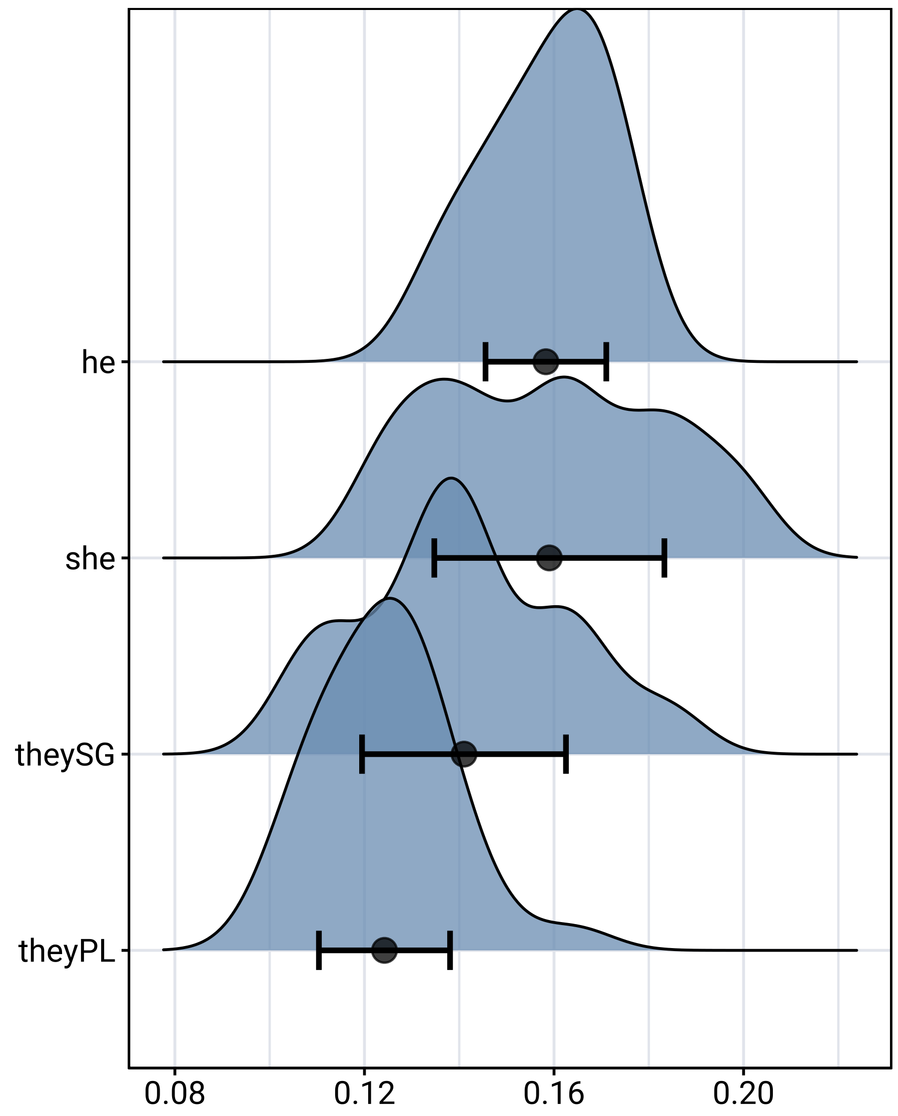

Semantic measures

SEMANTIC NEIGHBOURHOOD DENSITY

Correlation of a target with its

20 nearest neighbours

Semantic measures

SEMANTIC NEIGHBOURHOOD DENSITY

Correlation of a target with its

20 nearest neighbours

- Higher values indicate more semantically close neighbours

Semantic measures

SEMANTIC NEIGHBOURHOOD DENSITY

Correlation of a target with its

20 nearest neighbours

- Higher values indicate more semantically close neighbours

singular *they*

↓

plural *they*

↓

he, she

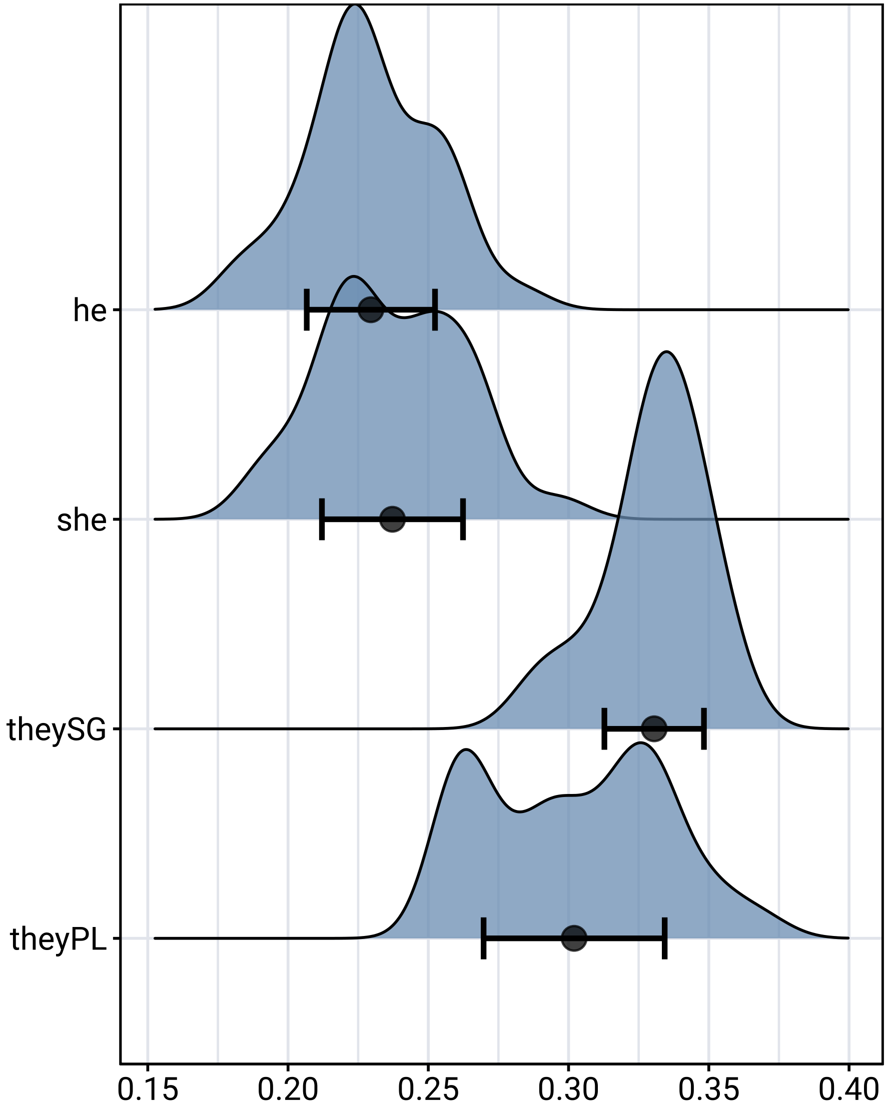

Discussion

Research questions

Research questions

RQ1 – Methodological Question

Can distributional semantics meaningfully capture pronoun semantics?

Research questions

RQ1 – Methodological Question

Can distributional semantics meaningfully capture pronoun semantics?

→ **YES**

Research questions

RQ1 – Methodological Question

Can distributional semantics meaningfully capture pronoun semantics?

→ **YES**

RQ2 – Theoretical Question

How is singular *they* semantically related to other third-person pronouns?

Research questions

RQ1 – Methodological Question

Can distributional semantics meaningfully capture pronoun semantics?

→ **YES**

RQ2 – Theoretical Question

How is singular *they* semantically related to other third-person pronouns?

→ **well…**

Semantic measures

Semantic measures

- Semantic measures
  - derived from an LDL implementation
  - making use of NDL and instance vectors

Semantic measures

- Semantic measures
  - derived from an LDL implementation
  - making use of NDL and instance vectors

capture the relation of third-person pronouns

Semantic measures

- Semantic measures
  - derived from an LDL implementation
  - making use of NDL and instance vectors

  capture the relation of third-person pronouns

- **SEMANTIC ACTIVATION DIVERSITY**
  - singular *they* in-between *he*, *she* and plural *they*

    = situated between its singular competitors and its plural homophone

Semantic measures

- Semantic measures
  - derived from an LDL implementation
  - making use of NDL and instance vectors

capture the relation of third-person pronouns

- **SEMANTIC ACTIVATION DIVERSITY**
  - singular *they* in-between *he*, *she* and plural *they*

    = situated between its singular competitors and its plural homophone
- **SEMANTIC NEIGHBOURHOOD DENSITY**
  - singular *they* has highest neighbourhood density

    = potential effect of belonging to two “worlds” – singular and plural pronouns

Conclusion

Conclusion

- Singular *they* semantics are situated somewhere between other third-person singular pronouns and plural *they*

Conclusion

- Singular *they* semantics are situated somewhere between other third-person singular pronouns and plural *they*
- This hints at previous research stating that singular *they* – a singular pronoun – retains parts of the plurality of plural *they* (Sanford & Filik, 2007)

Conclusion

- Singular *they* semantics are situated somewhere between other third-person singular pronouns and plural *they*
- This hints at previous research stating that singular *they* – a singular pronoun – retains parts of the plurality of plural *they* (Sanford & Filik, 2007)

Thank you!

References

Baayen, R. H., Chuang, Y.-Y., Shafaei-Bajestan, E., & Blevins, J. P. (2019). The discriminative lexicon: A unified computational model for the lexicon and lexical processing in comprehension and production grounded not in (de)composition but in linear discriminative learning. *Complexity*, 2019, 4895891. https://doi.org/10.1155/2019/4895891

Baayen, R. H., Milin, P., Đurđević, D. F., Hendrix, P., & Marelli, M. (2011). An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. *Psychological Review, 118*(3), 438–481. https://doi.org/10.1037/a0023851

Chuang, Y.-Y., Lõo, K., Blevins, J. P., & Baayen, R. H. (2020). Estonian case inflection made simple: A case study in Word and Paradigm Morphology with Linear Discriminative Learning. In L. Körtvélyessy & P. Štekauer (Eds.), *Complex words* (pp. 119–141). Cambridge University Press.

Conrod, K. (2020). Pronouns and gender in language. In *The Oxford Handbook of Language and Sexuality*. Oxford University Press. https://doi.org/10.1093/oxfordhb/9780190212926.013.63

Davies, M. (2008). *The Corpus of Contemporary American English (COCA)*. https://www.english-corpora.org/coca/

Han, C. H., & Moulton, K. (2022). Processing bound-variable singular they. *Canadian Journal of Linguistics/Revue Canadienne de Linguistique, 67*(3), 267–301. https://doi.org/10.1017/CNJ.2022.30

Harris, Z. S. (1954). Distributional structure. *WORD, 10*(2–3), 146–162. https://doi.org/10.1080/00437956.1954.11659520

Konnelly, L., Cowper, E., Konnelly, L., & Cowper, E. (2020). Gender diversity and morphosyntax: An account of singular they. *Glossa: A Journal of General Linguistics, 5*(1). https://doi.org/10.5334/GJGL.1000

Lapesa, G., Kawaletz, L., Plag, I., Andreou, M., Kisselew, M., & Padó, S. (2018). Disambiguation of newly derived nominalizations in context: A Distributional Semantics approach. *Word Structure, 11*(3), 277–312. https://doi.org/10.3366/word.2018.0131

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). Appleton-Century-Crofts.

Sanford, A. J., & Filik, R. (2007). “They” as a gender-unspecified singular pronoun: Eye tracking reveals a processing cost. *Quarterly Journal of Experimental Psychology, 60*(2), 171–178. https://doi.org/10.1080/17470210600973390

Schmid, H. (1999). Improvements in part-of-speech tagging with an application to German. In S. Armstrong, K. Church, P. Isabelle, S. Manzi, E. Tzoukermann, & D. Yarowsky (Eds.), *Natural language processing using very large corpora* (pp. 13–25). Springer. https://doi.org/10.1007/978-94-017-2390-9_2

Schmitz, D., Plag, I., Baer-Henney, D., & Stein, S. D. (2021). Durational differences of word-final /s/ emerge from the lexicon: Modelling morpho-phonetic effects in pseudowords with linear discriminative learning. *Frontiers in Psychology*, 12. https://doi.org/10.3389/fpsyg.2021.680889

Schmitz, D., Schneider, V., & Esser, J. (2023). No genericity in sight: An exploration of the semantics of masculine generics in German. *Glossa Psycholinguistics*. Preprint available on PsyArXiv. doi: 10.31234/osf.io/c27r9

Wagner, A. R., & Rescorla, R. A. (1972). Inhibition in Pavlovian conditioning: Application of a theory. In R. A. Boakes & M. S. Halliday (Eds.), *Inhibition and learning* (pp. 301–334). Academic Press Inc.

Semantic measures

Semantic measures

- Random forest analyses show that

Semantic measures

- Random forest analyses show that
  - regarding **SEMANTIC ACTIVATION DIVERSITY**,

    singular *they* is most frequently confused with *he*, *she*, and plural *they*

Semantic measures

- Random forest analyses show that
  - regarding **SEMANTIC ACTIVATION DIVERSITY**,

    singular *they* is most frequently confused with *he*, *she*, and plural *they*
  - regarding **SEMANTIC NEIGHBOURHOOD DENSITY**,

    singular *they* is most frequently confused with *anybody*, *anyone*, and plural *they*

Semantic space

Semantic space