Eight shades of animacy

Property axes and split intransitivity

Eva Neu, Brian Dillon and Katrin Erk

PLC 50 | February 28, 2026

University of Massachusetts Amherst

The big picture

The lexical meaning of verbs affects their syntactic behavior:

(1) a. I sent you a book.

b. *I slept him the pillow.

The big picture

Our focus here: split intransitivity

(2) a. the broken glass

b. *the played girl

The big picture

Two general questions:

- Empirically, which semantic properties matter for the syntactic behavior of verbs?
- Theoretically, how is this knowledge about the relation between lexical meaning and syntax encoded?

The big picture

Today's contribution:

- A new perspective: let's approach these questions using word embeddings!
- A new methodology for working with embeddings: rating-based property axes
- Some first results

Outline

1. Background: Split intransitivity
2. Methodology
3. Modeling and results
4. Conclusion

1. Background: Split intransitivity

Split intransitivity

Two kinds of intransitives (Burzio, 1981, 1986; Perlmutter, 1978):

- Unergatives: sole argument is syntactically and semantically identical to the subject argument of transitives
  - External argument position
  - Agent interpretation
- Unaccusatives: sole argument is syntactically and semantically identical to the object argument of transitives
  - Internal argument position
  - Patient interpretation

Split intransitivity

(3) Unergative

a. The girl played.

b.

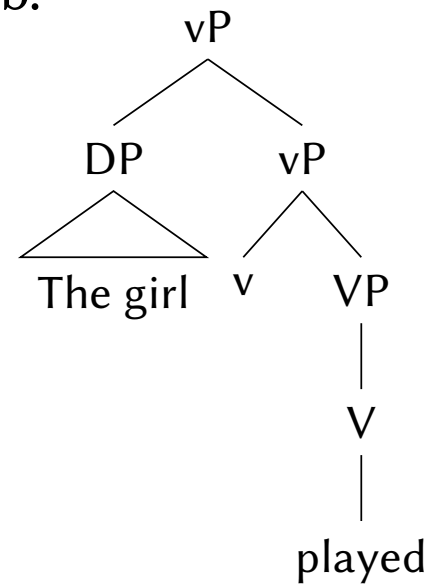

(4) Unaccusative

a. The glass broke.

b.

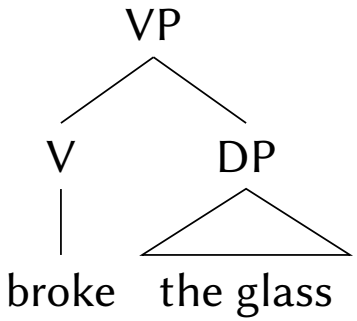

Split intransitivity

Unaccusativity diagnostics such as reduced relatives detect a difference between the two structures:

(5) a. the broken glass — unaccusative

b. *the played girl — unergative

Other diagnostics: resultative secondary predicates, impersonal passives, auxiliary selection, agent nominalizations, ...

Split intransitivity

Verbs can allow for both an unergative and an unaccusative structure...

(6) a. *uṛ-ii (huu-ii) ciṛyaa
fly-PFV.FSG be-PFV.FSG bird.FSG
Intended: 'the flown bird'

b. uṛ-ii (huu-ii) patang
fly-PFV.FSG be-PFV.FSG kite.FSG
'the flown kite' (Ahmed, 2010:8f.)

... but it is not the case that anything goes.

Split intransitivity

How is the syntactic behavior of intransitives constrained?

- Verbs have a gradient tendency towards one structure or the other depending on their semantics (Sorace, 2000, 2004, 2011)
- Can we confirm which semantic properties matter? And how is this gradience encoded?

2. Methodology

Word embeddings

What are word embeddings?

- Numerical representations of word meaning in vector space
- GloVe: model trained on word co-occurrence statistics
  $\rightarrow$ words that occur in similar contexts have similar vector representations

Word embeddings

Leveraging word embeddings for the problem of encoding split intransitivity:

- Identify a dimension in vector space that corresponds to Sorace's unergative-unaccusative spectrum
- Determine which semantic properties are associated with this dimension using property axes

Property axes

Property axes are dimensions in embedding space that encode gradable properties (e.g., Grand et al., 2022; Kozlowski et al., 2019):

- Danger
- Size
- Wealth
- ...

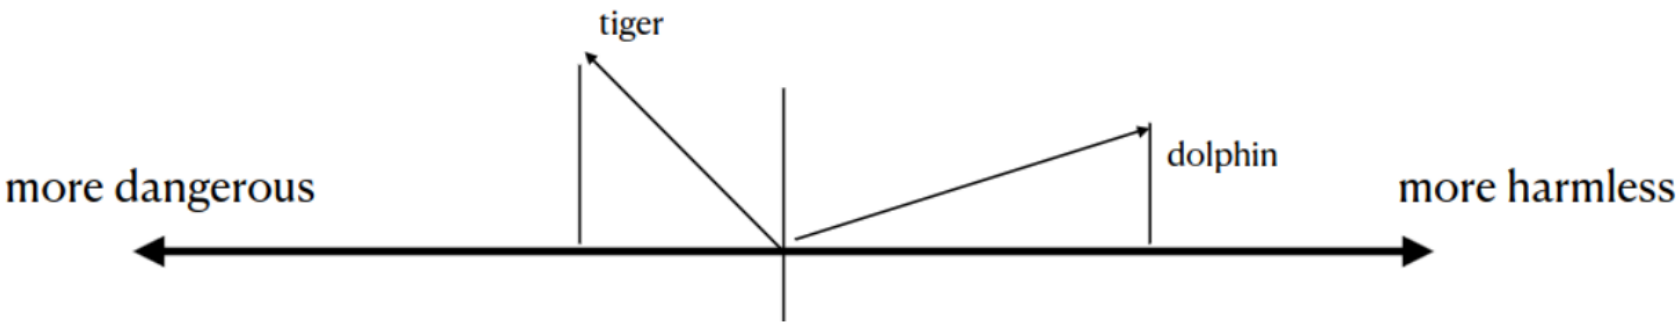

Property axes

A new method for computing property axes: rating-based axes

- Collect human ratings for a particular semantic property
- Project an axis in space that optimizes the fit to these ratings

  (see also Erk and Apidianaki, 2024)

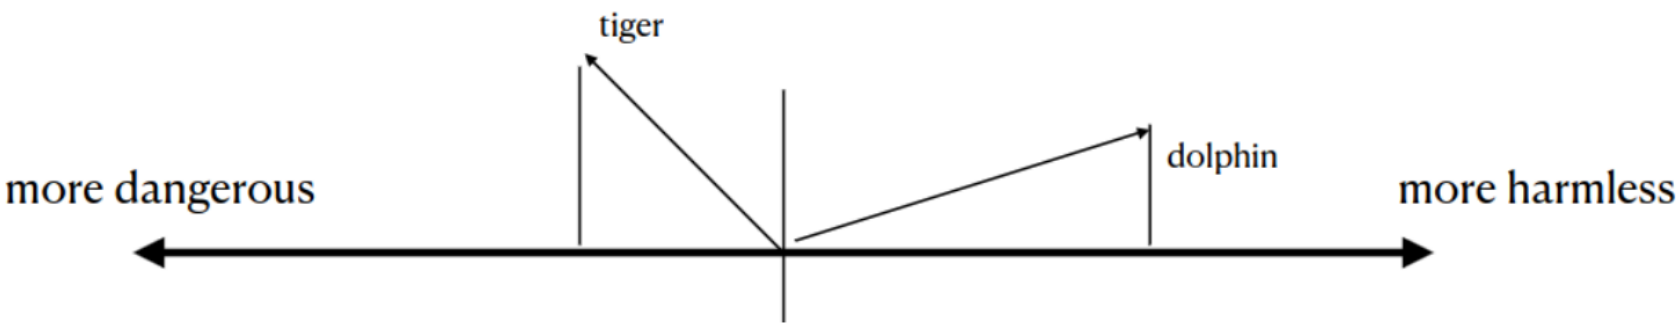

3. Modeling and results

Modeling

Our strategy:

- Focus on one semantic property: agentivity (Dowty, 1991)
- Identify axes in embedding space that encode various flavours of agentivity
- Measure how high different verbs score on these properties
- Compute how well these measures predict unergative/unaccusative behavior

Modeling

Our syntactic data: acceptability ratings for 138 intransitive verbs in reduced relative clauses, an unaccusativity diagnostic (Kim et al., 2024)

(7) a. the frozen ground

b. the abounded opportunities

c. the hopped frogs

d. ...

Modeling

Our semantic data: ratings for 1,200 concrete nouns on 6 different animacy dimensions (VanArsdall and Blunt, 2022)

- General living/non-living scale
- Ability to think
- Ability to reproduce
- Similarity to a person
- Goal-directedness
- Movement likelihood

Further clustered via factor analysis into two broader dimensions:

- Mental animacy
- Physical animacy

Modeling

Step-by-step procedure:

- For each of VanArsdall and Blunt's animacy dimensions, compute a property axis that optimizes the fit to the human ratings
- Cross-validate the axis on the original data set
- Project Kim et al.'s 138 intransitives onto these axes to derive a semantic measure
- Use this semantic measure to predict Kim et al.'s syntactic ratings in a Bayesian regression model (`brm`)
- Compare the performance of the different regression models (LOO analysis)

Results

Models	ELPD diff	SD
Move vs. null model	444.6	28.5
Living vs. null model	214.4	2.6
Thought vs. null model	195.5	19.2
Person vs. null model	184.3	18.8
Reproduction vs. null model	149.2	17.2
Goals vs. null model	74.3	12.1
Mental vs. null model	156.3	17.4
Physical vs. null model	23.5	21.2

Table 1: LOO analysis with type-level animacy axes

(ELPD = expected log predictive density)

Results

- Caveat: word type embeddings conflate transitive and intransitive verbs
- One possible solution: use a language with unambiguous morphological marking
- Another option: token embeddings – ask me in the Q&A!

4. Conclusion

Conclusion

Recall our two big questions from the beginning:

- Empirically, which semantic properties matter for the syntactic behavior of verbs?
    - Animacy in the broad sense (moving, being alive) is a better predictor of the syntactic behavior of intransitives than agentivity in the narrow sense (having goals, being a person)
- Theoretically, how is this knowledge about the relation between lexical meaning and syntax encoded?
    - Our approach: the syntactic behavior of verbs is a function of the embedding space

Conclusion

Future directions:

- Use token-level instead of type-level axes
- Move beyond a single unaccusativity diagnostic
- Compare property axes across languages

Acknowledgments

Thanks to Songhee Kim!

And thanks to you!

References

Ahmed, T. (2010). The unaccusativity/unergativity distinction in Urdu. Journal of South Asian Linguistics, 3, 3–22.

Burzio, L. (1981). Intranitive verbs and Italian auxiliaries [Doctoral dissertation, MIT].

Burzio, L. (1986). Italian syntax: A Government and Binding approach. D. Reidel.

Dowty, D. (1991). Thematic proto-roles and argument selection. Language, 67, 547–619.

Erk, K., & Apidianaki, M. (2024). Adjusting interpretable dimensions in embedding space with human judgments.

Grand, G., Blank, I. A., Pereira, F., & Fedorenko, E. (2022). Semantic projection recovers rich human knowledge of multiple object features from word embeddings. Nature Human Behaviour, 6, 975–987.

References

Kim, S., Binder, J. R., Humphries, C., & Conant, L. L. (2024). Decomposing unaccusativity: A statistical modelling approach. Language, Cognition and Neuroscience, 39, 1189–1211.

Kozlowski, A. C., Taddy, M., & Evans, J. A. (2019). The geometry of culture: Analyzing the meanings of class through word embeddings. American Sociological Review, 84, 905–949.

Levin, B., & Rappaport Hovav, M. (1995). Unaccusativity. At the syntax-lexical semantics interface. MIT Press.

Levin, B., & Rappaport Hovav, M. (2005). Argument realization. Cambridge UP.

Perlmutter, D. (1978). Impersonal passives and the Unaccusative Hypothesis. Papers from the Annual Meeting of the Berkeley Linguistic Society, 4, 157–189.

Sorace, A. (2000). Gradients in auxiliary selection with intransitive verbs. Language, 76, 859–890.

References

Sorace, A. (2004). Gradience at the lexicon-syntax interface: Evidence from auxiliary selection and implications for unaccusativity. In A. Alexiadou, E. Anagnostopoulou, & M. Everaert (Eds.), The unaccusativity puzzle (pp. 243–268). Oxford UP.

Sorace, A. (2011). Gradience in split intransitivity: The end of the Unaccusative Hypothesis? Archivio Glottologico Italiano, 96, 67–86.

VanArsdall, J. E., & Blunt, J. R. (2022). Analyzing the structure of animacy: Exploring relationships among six new animacy and 15 existing normative dimensions for 1,200 concrete nouns. Memory & Cognition, 50, 997–1012.

Appendix

Models	ELPD diff	SD
Move vs. null model	305.8	24.4
Living vs. null model	214.5	20.6
Thought vs. null model	178.0	18.7
Reproduction vs. null model	177.5	18.7
Goals vs. null model	101.6	14.2
Person vs. null model	95.7	13.9

Table 2: LOO analysis with token-level animacy axes