Evidence from property axes on the semantic correlates of split intransitivity

Eva Neu, Brian Dillon and Katrin Erk

CLS 62 | April 17–19, 2026

University of Massachusetts Amherst

Properties

Gradable properties are everywhere:

- How dangerous are different animals?
- How spicy are different dishes?
- How hungry are different people in the audience?
- How interesting are different conference talks?

Properties

How can we quantify gradable properties?

- The traditional strategy: human ratings
- A computational approach: property axes in embedding space

Properties

Today's talk:

- Develop a new methodology for detecting property axes in embedding space
- Apply this methodology to an old problem in argument structure: the effect of gradable semantic properties on split intransitivity

Roadmap

- Background
- Methods
- Results
- Discussion

Background

Two kinds of intransitives:

- 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

(Burzio, 1981, 1986; Perlmutter, 1978)

Background

(1) Unergative

a. The girl played.

b.

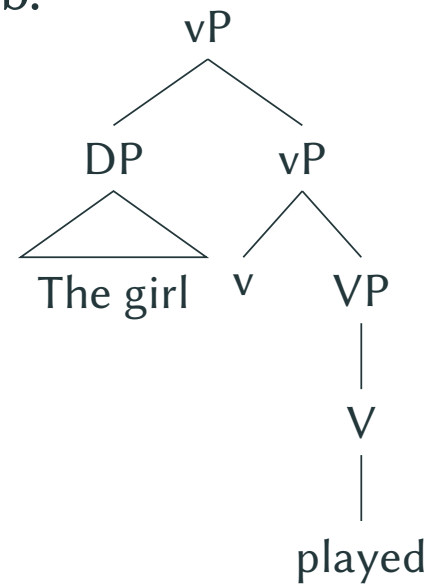

(2) Unaccusative

a. The glass broke.

b.

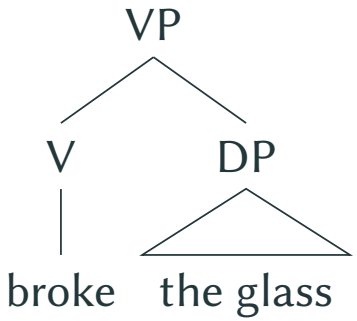

Background

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

(3) a. the broken glass — unaccusative

b. *the played girl — unergative

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

Background

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

(4) 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. Verbs have a gradient tendency towards one structure or the other.

Background

It has been argued that the syntactic behavior of a verb is affected by gradable semantic properties:

- The more *agentive* the meaning of a verb is, the more likely it is to surface in an *unergative* syntax,
- The more *telic* the meaning of a verb is, the more likely it is to surface in an *unaccusative* syntax.

However, robust empirical support for these and similar claims is limited.

(Sorace, 2000, 2004, 2011, a.o.)

Background

Our goal today is not to solve the general question of which semantic properties affect split intransitivity. Rather:

- We focus on one semantic property that appears to be the most promising candidate for predicting split intransitivity: agentivity.
- We aim to make more precise which flavour of agentivity syntax is sensitive to.
- In doing so, we develop a methodology that can then be used to tackle this question more broadly, as well as a range of other problems.

Roadmap

- Background
- Methods
- Results
- Discussion

Methods

What are word embeddings?

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

Methods

Property axes are dimensions in embedding space that encode gradable properties:

- Danger
- Size
- Wealth
- ...

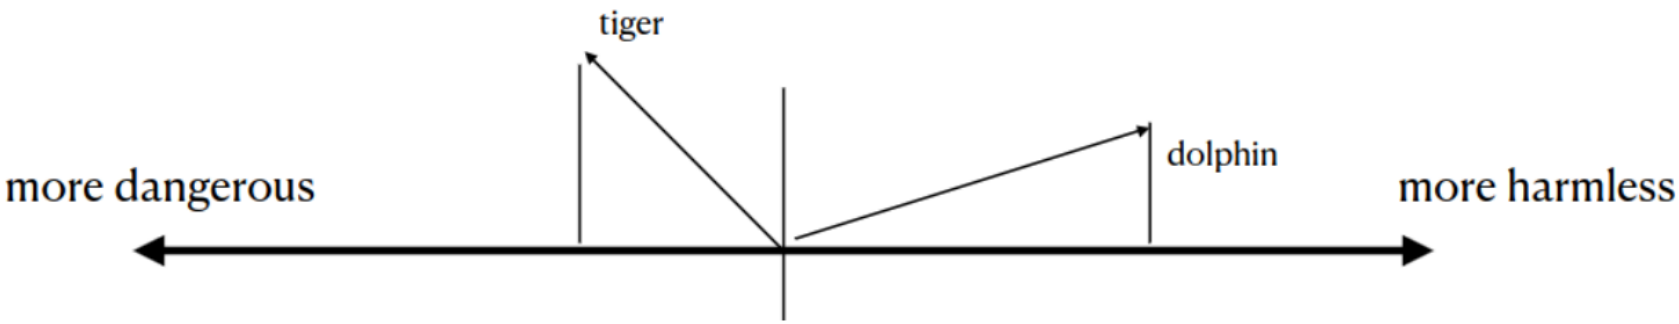

(e.g., Grand et al., 2022; Kozlowski et al., 2019)

Methods

We develop a new method for computing property axes: rating-based axes.

- Collect human ratings for a particular semantic property
- Find an axis in space that optimizes the fit to these ratings using a ranking loss function
- Extrapolate measures for unseen words by projecting their embeddings onto this axis

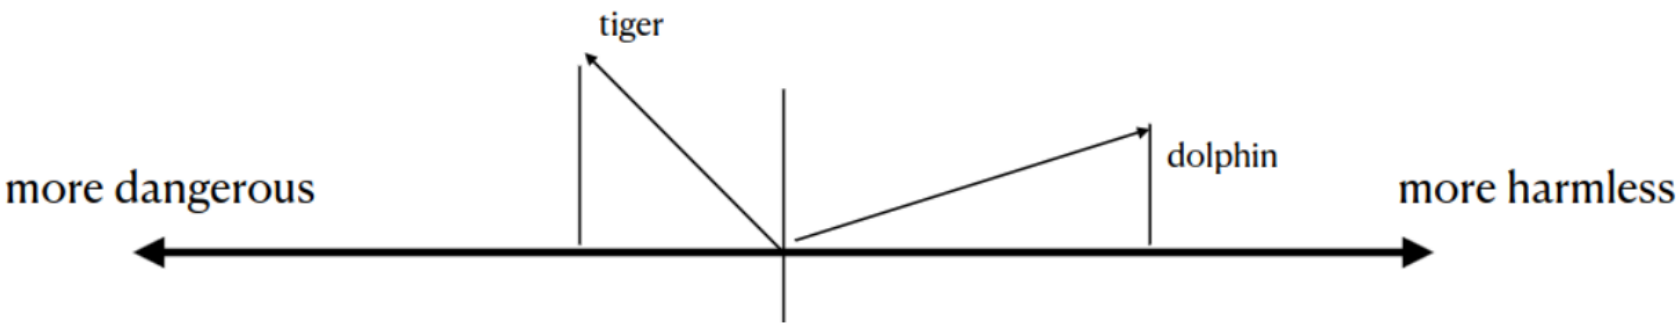

(see also Erk and Apidianaki, 2024)

Methods

How we use rating-based axes for the problem of split intransitivity:

- Identify axes in embedding space that encode various flavours of agentivity based on ratings for nouns
- Measure how high different verbs score on these properties by projecting them on the axes
- Compute how well these measures predict unergative/unaccusative behavior

Methods

Our syntactic data: acceptability ratings for 138 intransitive verbs in reduced relative clauses, an unaccusativity diagnostic

(5) a. the frozen ground

b. the abounded opportunities

c. the hopped frogs

d. ...

(Kim et al., 2024)

Methods

Our semantic data: ratings for 1,200 concrete nouns on 6 different animacy dimensions

- 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

(VanArsdall and Blunt, 2022)

Methods

The 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 and other data (details in the appendix)
- Project Kim et al.'s 138 intransitives onto these axes to derive a semantic measure
- Test how well these different semantic measures predict the syntactic ratings in a Bayesian `brms` model
- Compare the performance of the different predictors with an LOO analysis

Roadmap

- Background
- Methods
- Results
- Discussion

Results

Estimates for fixed-effect coefficients with estimate error and upper and lower bound 95% confidence interval:

Predictor	Est.	Est. error	l-95% CI	u-95% I
Living	-0.24	0.01	-0.26	-0.21
Thought	-0.23	0.01	-0.25	-0.20
Repr.	-0.20	0.01	-0.22	-0.18
Person	-0.22	0.01	-0.24	-0.20
Goals	-0.14	0.01	-0.16	-0.12
Move	-0.35	0.01	-0.37	-0.32
Mental	-0.20	0.01	-0.23	-0.18
Physical	-0.25	0.01	-0.27	-0.22

All fixed effects are negative: more animate verbs are less acceptable in reduced relative clauses.

Results

Results from LOO analysis with type-level animacy axes; ELPD = expected log-predictive density:

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

Results

- Animacy in the broad sense (moving, being alive) is a better predictor than a narrow notion of goal-directedness or intentionality.
- Why does Physical nonetheless perform worse than Mental?
    - Does not account for much variance in VanArsdall and Blunt's dataset in the first place
    - Not strongly correlated with Movement

Roadmap

- Background
- Methods
- Results
- Discussion

Discussion

Some limitations of this study:

- Focused on a single semantic dimension, a single unaccusativity diagnostic and a single language
- Used only a single morphological form for the verb
- Did not specifically target intransitives

Discussion

Moving forward:

- Use token-level embeddings that allow us to specifically extract intransitive uses of a verb
- Confirm more directly that ratings collected for nouns carry over to other parts of speech
- Turn the problem on its head: let's find axes that predict the syntactic data well and then determine which semantic property it encodes

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: cross-validation results

We cross-validated the rating-based dimensions fitted with ranking loss on the Grand et al. (2022) dataset and compared them to two other sets of property axes:

- Seed-based axes (Grand et al., 2022)
- Seed-based axes interpolated with human ratings, fitted with pointwise loss (Erk and Apidianiki, 2024)

Model	POC	XPOC
Seed	.629	.631
Pointwise	.701	.779
Ranking	.703	.797

Appendix: cross-validation results

We further compared the performance of Erk and Apidianiki's pointwise loss model without seeds and the new ranking loss model on the VanArsdall and Blunt (2022) data:

Model	POC	pearson *r*
Pointwise	.516	.05
Ranking	.783	.789

Appendix: cross-validation results

Lastly, we computed the correlation of the eight animacy axes to ratings collected by Kim et al. (2024) on how intentionally the action described by the verb was performed:

- Significant correlation for Goals (.286, *p*-val. .0), Thought (.191, *p*-val. .001), Person (.191, *p*-val. .001), and coarse-grained Mental Animacy (.254, *p*-val. .0)
- No correlation for the other features