

# Disentangling factors in the placement of manner adverbials in German: The effect of distributional similarity

Martin Schäfer (post@martinschaefer.info)

SFB 833: The Construction of Meaning, Universität Tübingen, Germany

Holger Gauza (holger.gauza@uni-tuebingen.de)

SFB 833: The Construction of Meaning, Universität Tübingen, Germany

Britta Stolterfoht (britta.stolterfoht@uni-tuebingen.de)

German Department, Universität Tübingen, Germany

## Abstract

Processing differences between obligatory constituents like verbal complements and facultative ones like adjuncts are widely discussed in psycholinguistics. But the relative positional and semantic variability of adjuncts makes analyzing these differences a daunting challenge. Focusing on the intricate problem of the default position of manner adverbials in German, we present a re-analysis of a recent psycholinguistic study on their ordering preferences in which we explore the extent to which similarities between word embeddings can be used as stand-ins for shared semantic memory, representing the probability of seamless conceptual combination. In re-analyzing six experiments across different paradigms, the addition of the new predictors yields substantially better models which show that these factors considerably interact with established lexical and grammatical predictors.

**Keywords:** language processing; word order; argument; adjunct; conceptual combination; distributional semantics; verb; direct object; manner adverbial; German

## Introduction

Adequately understanding the processing of adjuncts as opposed to that of syntactic arguments is complicated by many factors. Manner adverbials in German are a case in point. One core challenge concerns their placement relative to the arguments of the verb, in particular the verb's direct object: Which order, (1-a), **object** > **adverbial**, or (1-b), **adverbial** > **object**, shows them in their default position?

- (1) a. [...], weil Elin eine Tür schnell öffnet.  
[...], because Elin a door quick opens  
b. [...], weil Elin schnell eine Tür öffnet.  
[...], because Elin quick a door opens  
“because Elin (quickly) opens a door (quickly)”

Prominent factors include the status of the direct object in terms of information structure as well as in relation to the verb, that is, whether the direct object is integrated or not (Frey & Pittner, 1998). The possible readings of the adverbial itself are a further important factor, see Schäfer (2013) and Gauza (2018) for comprehensive overviews. This study proposes a rethinking of the concept of integration proposed by Jacobs (1993) in terms of shared semantic memory, and explores this approach via new analyses of the data from a slew of experiments from Gauza (2018), using the distributional similarity between verb, direct object, and manner adverbial as stand-in for a language user's experience.

Integration is introduced in Jacobs (1993) for two constituents constituting a single referential unit and requiring only one step of semantic processing. He exemplifies this for verbs and objects with the pair from example (1), “eine Tür öffnen” ‘to open a door’. According to him, this is an instance of integration, because its two constituents form a compact semantic unit which establishes reference in a single step: the phrase is not relating a specific door to a specific door-opening event, but establishes a holistic reference to door-opening events. With respect to the default position of manner adverbials relative to the direct object, this concept is taken up in Frey and Pittner (1998), who assume a default order **object** > **adverbial** but argue that the reverse order becomes the default in cases of object integration. Consider example (1) again: by default, their account predicts the ordering in (1-a). But because they assume that object integration takes preference over the expected default order, the default order in this case changes to **adverbial** > **object**, as in (1-b).

For arguments, most conditions for integration concern morpho-syntactic constraints (eg. c-command), but there are also information structural constraints. Eckardt (2003) postulates a clear difference between indefinite and definite direct objects: in sentences with manner adverbials, definites need to be topical, and are moved to a position before the adverbial. Not adjacent to the verb anymore, they cannot be integrated. Semantic conditions for integration are left vague (eg. prototypical object properties).

In contrast, this study is interested in a quantitative exploration of the semantic preconditions for integration. We see integration as a special case of conceptual combination, hypothesizing that in order to allow for integration, two constituents must be tightly linked in the speaker's experience, making conceptual combination completely seamless. On this view, *to open a door* is a prime candidate for integration because both concepts are experientially linked due to the very function of door itself, as well as the fact that the event itself constitutes a common everyday experience. In contrast, *to open a grave*, although the direct object is also an indefinite noun phrase, is not a prime candidate for integration because no or only weak experiential links exist between the two concepts. In terms of verbal arguments, this approach adopts the view of thematic roles as verb-specific concepts, incorporating experiential knowledge of events and

their typical participants (McRae, Ferretti, & Amyote, 1997).

This focus on the semantic preconditions allows us to address an issue often overlooked in work on object integration and word order: Jacobs (1993) assumes that modifiers can also be integrated. Viewing integration as the seamless combination of two experientially linked concepts allows to quantify the cognitive precondition for integration for both manner adverbials and direct objects alike. While conceptually manner adverbials and direct objects typically differ in that the former prototypically encode property concepts and the latter object concepts, this intuitive idea of integration can be applied to both. For example, integration is expected for combinations of verbs of sound emission and adverbials modifying sound emission (eg. *laut singen* ‘loudly sing’), and similarly for verbs of motion and speed adverbials, e.g. *schnell laufen* ‘quickly run’. If this is on the right track, we expect an influence of this aspect not only on the integrability of two constituents, but we also expect competition between adverbials and direct objects that both allow for seamless integration with the verb, e.g. *laut singen* ‘to sing loudly’ and *ein Lied singen* ‘to sing a song’. Note that this view of integration leaves open the internal process of conceptual combination used. That is, we are not disputing that there are many different ways in which conceptual combination may proceed, nor that internal processes for the combination of a thing-concept and an event-concept might be very different from those required for the combination of a property-concept and an event-concept. Rather, we are simply assuming that several types of concepts can be equally seamlessly combined, and it is this ease of combination that our measure of shared experience is targeting.

As a stand-in for the speakers experience and expectation with regards to the combination of two concepts, we use the cosine similarity between the distributional semantic representations of the constituents involved. With Günther, Rinaldi, and Marelli (2019), we assume that distributional semantic models, that is, high-dimensional semantic spaces, can be meaningfully treated as general-level models of semantic memory, encoding the speakers experience with specific contexts (Cf. also Lebani and Lenci’s 2018 distributional approach to thematic role concepts).

The semantic similarity between two distributional semantic representations can consequently be seen as the degree of shared experience the speaker associates with the two constituents, and with that as a measure of seamlessness of integration in case of noncompeting concepts, or as a measure of competition for possibly competing concepts. The three cosine similarities of interest for the placement of manner adverbials relative to direct objects are the adverbial-verb similarity, the object-verb similarity, and the object-adverbial similarity. For the first two, adverbial-verb and object-verb similarity, the cosine similarity can be seen as gauging the potential of seamless integration. The object-adverbial similarity holds between concepts that are not candidates for integration because objects and adverbials are necessarily linked via

the event provided by the verb. Because direct objects and adverbials compete for the same linear positions in the sentences of interest here, their similarity can be understood as a measure of competition.

In remodeling experiments from Gauza (2018), we explore whether these three new predictors lead to better models or significantly interact with the predictors already considered in that study, with the main focus on their effects with regard to word order.

## Experimental data

We used the raw data from Gauza (2018). To avoid cherry-picking, we re-modeled the complete first set of experiments reported there, which aimed to establish the default position of manner adverbials relative to the direct object.

The experiments include manipulations of core structural and lexical factors hypothesized to influence the positioning of manner adverbials. The order varies between **adverbial > object** and **object > adverbial** (factor **order**). The direct objects are either definite vs. indefinite nominals (factor **object status**), which allows to probe the interaction of information structure with word order. The manner adverbials used are exclusively adjectives, varying in the factor **inchoative**. The set of +inchoative adjectives are speed adjectives like *schnell* ‘quickly’ in (1), which allow non-manner readings: either inchoative readings indicating a short time from the onset of some contextually given point to the onset of the event described by the verbal predicate, or holistic readings, indicating that the whole event took only a short time. Both readings are associated with a higher syntactic attachment site corresponding to a linear position before the direct object (Schäfer, 2013), comparable to their behavior in English (Ernst, 2002).

The status of the direct object was also manipulated across experiments, being either contextually available or not. Experiments III, IVa, and IVb used patterns argued to be syntactic test environments in Frey and Pittner (1998): in Experiment III, either [adverbial + verb] or [direct object + verb] were topicalized, in Experiments IVa and IVb, only the indefinite *w*-pronoun *was* ‘what’ served as the direct object. Across different experimental paradigms, the material used stayed largely the same. Two sentences from Experiment Ia illustrate the different factor levels:

- (2) a. Elisabeth sagt, dass Björn das Gedicht laut  
Elisabeth says that Björn the poem loud  
vorgetragen hat.  
recited has  
‘Elisabeth says that Björn recited the poem loudly.’
- b. Susanne erzählt, dass Sybille schnell einen  
Susanne reports that Sybille quick a  
Roman gelesen hat.  
novel read has  
‘Susanne reports that Sybille quickly read a novel.’

In (2-a), the direct object is definite and precedes the - inchoative adjective. In (2-b), the +inchoative adjective precedes the indefinite object. Table 1 provides an overview of the experiments.

Table 1: Overview of first set of experiments in Gauza (2018)

Exp.	Method	Format
Ia	Likert scale	subordinate clause
Ib	Fragment arrangement	same as Ia
II	Likert scale	= Ia but given direct object
III	Likert scale	ADV/DO + verb preposed
IVa	Likert scale	=Ia but direct object indefinite w-pronoun)
IVb	Self-paced reading	same as IVa

Experiment Ia, in which **inchoative** was a post-hoc factor, also contained non-manner adverbials. These are completely excluded in our reanalysis.

## Hypotheses

We expect that the two verb-related similarity measures interact with the factors inchoative and object status, respectively.

A. For -inchoative adverbials, the manner interpretations are the only ones available and since they are directly associated with the specifics of the verbal predicate, we expect a constant influence of adverbial-verb similarity across syntactic positions: the higher the similarity, the higher the likelihood of integration, the better the position after the object. For +inchoative adverbials in their holistic and inchoative readings, the specifics of the verbal predicate are hypothesized to be less important since they are not manner adverbials. We therefore expect a muted influence of the adverbial-verb similarities on these readings, which are in turn associated with a higher syntactic position.

B. If definite noun phrases already prohibit integration because they require syntactic movement (Eckardt, 2003), then the object-verb similarity is irrelevant in this condition and should therefore have no influence on word order. For indefinites, the greater the similarity, the higher the likelihood of integration, and therefore the greater the preference for verb-adjacent position.

C. The object-adverbial similarity is hypothesized to be a measure of competition: if both direct object and manner adverbial are associated with similar experiences, they might compete for the same slots. If this is the case, interaction with the factor **inchoative** is expected: only +inchoatives include the possibility of choosing an alternative reading.

## Methods

### Distributional predictors

The distributional semantics model was trained on the German Wikipedia using the word2vec algorithm

(Mikolov, Chen, Corrado, & Dean, 2013). The pre-trained vectors used here are provided by deepset.ai <https://deepset.ai/german-word-embeddings>. The pretrained vectors were used to calculate the cosine similarities of the three new predictors: adverbial-verb similarity, object-verb similarity, and object-adverbial similarity.

### Data trimming and cleaning

Items that contained target constituents for which no pre-trained vector was available (due to them falling outside of the 100,000 most frequent words) were excluded. This concerns one item in all experiments (using the verb *umtopfen* “repot”), and one further item in experiments IVa/b (*dribbeln* “dribble”). In IVa/b, one further item was excluded because in the combination (*was*) *laut* (*was*) *hören* ‘listen to sth. loud’ *laut* does not function as a manner adverbial.

For the acceptability judgements, subjects were excluded if they only used one single scale item or two levels with one being used more than 85% of the time. This led to the exclusion of four subjects in Ia, two subjects in III, and five subjects in IVa, corresponding to 6.25%, 5.56% 10.41% of all data points. For the self-paced reading task in Experiment IVb, data points exceeding 2000 milliseconds on the critical segments were excluded and afterwards the data was trimmed to include only reading times falling within 2.5 standard deviations. This affected 130 datapoints (10.94% of the total).

Trimming and cleaning did not change the effect of the original predictors as reported in Gauza (2018).

### Statistical analysis

The same modeling strategy was used across all data sets: we used linear mixed effects regression models, with random intercepts for items and subjects. In a first step, we remodeled the data without the distributional predictors. Initial models allowed all predictors to interact. To remove non-significant interactions of predictors as well as non-significant single predictors, we used the `step()` function from the `lmerTest` package (Kuznetsova, Brockhoff, & Christensen, 2017) to perform automatic backward elimination on the fixed effects in the initial model (except for the logistic regression analysis of experiment Ib, where we performed manual backward selection). In a second step, this procedure was repeated including the distributional predictors. ANOVAs were used for model comparison. Models with distributional predictors were only considered if their AIC was lower and the residual sum of squares was significantly lower than that of the best model using non-distributional predictors only. Reported t-test values for factors used the Satterthwaite’s methods from the `lmerTest` package. To represent the variance explained by the fixed effects, we use the marginal pseudo-R-squared value calculated using the `MuMIn` package (Bartoń, 2016). Figures were produced with the `effects` package (Fox, 2003).

## Results

In all cases, the addition of the distributional predictors lead to significantly better models with lower AICs and, in all but

two cases, very substantial increases in the variance explained by the fixed effects alone, see the overview in Table 2.

Table 2: First set of experiments: model comparisons

Exp.	variance explained		model comparison	
	non-distrib.	distributional	$\chi^2/(t)$	p-value
Ia	4.09%	9.76%	85.332	< 0.001
Ib	12.11%	27.64%	38.989	< 0.001
II	0.63%	1.89%	10.637	= 0.014
III	4.76%	8.20%	39.425	< 0.001
IVa	2.14%	9.47%	54.971	< 0.001
IVb	[null model]	0.29%	4.387	= 0.036

Below, we report the results in detail for the first three experiments (Ia, Ib, and II) and summarily for the last three experiments (III, IVa, and IVb).

### Experiment Ia

The experimental data consisted of naturalness ratings on sentences following the pattern in Example (2). The original predictors show a significant three-way interaction between **position**, **object status**, and **inchoative**. For -inchoatives, definites in the **object** > **adverbial** order show clear facilitation, indefinites show no effect, while for +inchoatives post-object position also shows facilitation for definites, but lower ratings for indefinites.

In the final distributional model, all similarity predictors occur in three- and two-way interactions. Adverbial-verb similarity participates in a three-way interaction with **inchoative** and object-verb similarity ( $t = 2.912$ ,  $p < 0.01$ ), the only interaction for object-verb similarity, plotted in Figure 1. In

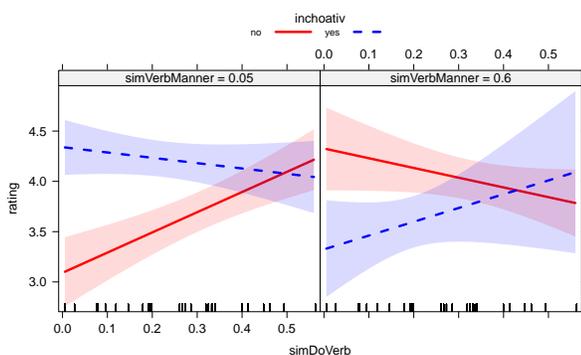


Figure 1: Experiment Ia, predictor effect plot for object-verb similarity and its interactions.

line with hypothesis A, the effect of adverbial-verb similarity is somewhat muted for the +inchoative items, compare the overall differences between the left and the right-hand panels. Notably, when the object-verb similarity is low, only the -inchoatives straightforwardly benefit from increased adverbial-verb similarity. Not in line with hypothesis

B, object-verb similarity does not interact with **object status**. Instead, for the -inchoatives the reversal in correlation between object-verb similarity and the ratings from clearly positive (left panel) to slightly negative might be indicative of competition between these two similarities.

Both other similarities interact with **position**. Adverbial-verb similarity via a two-way interaction ( $t = 4.456$ ,  $p < 0.001$ ), object-adverbial similarity via a three-way interaction modulated by **inchoative** ( $t = -5.461$ ,  $p < 0.001$ ). There continue to be interactions of **position** with both **object status** and **inchoative**. All interactions involving **position** are plotted in Figure 2. For the -inchoatives (left hand panels), both

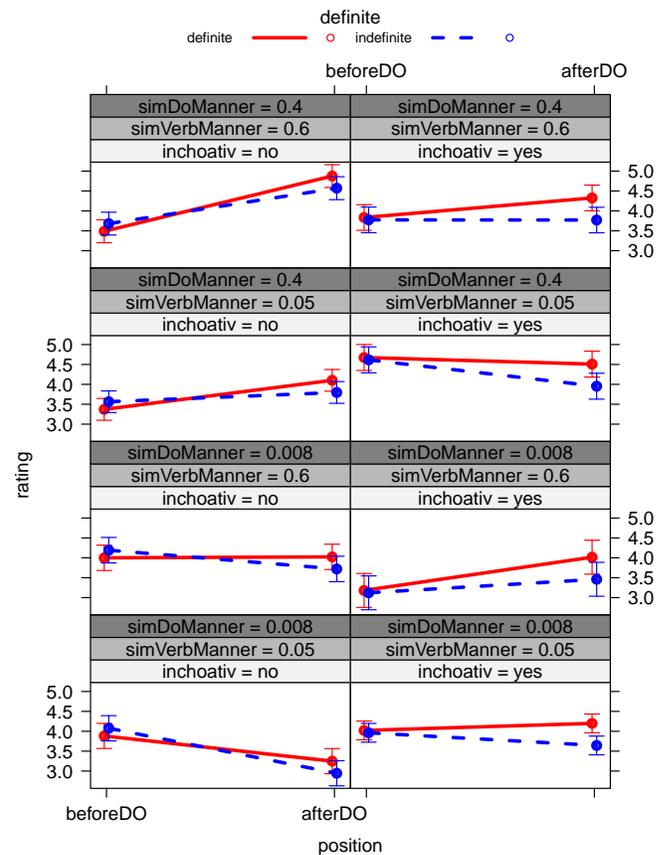


Figure 2: Experiment Ia, predictor effect plot for **position** and all its interactions.

indefinites and definites move in lockstep. When the similarity between object and manner adverbial is high, the order **object** > **adverbial** leads to better judgments, when it is low, they are lower. In conjunction with the adverbial-verb similarity, this leads to very clear positional effects when both values are high, where the **object** > **adverbial** order is clearly judged as more natural, whereas the reverse is true when both values are low. For the inchoatives, in the right hand panels, the model predicts consistently higher ratings for definites when the order is **object** > **adverbial**. However, across positions, an increase in ratings is only predicted when the adverbial-verb similarity is high. In the **adverbial** > **object**

order, indefinites and definites behave alike. However, in this condition we find very clear differences especially when one similarity is high and the other is low. When object-adverbial similarity is high but adverbial-verb similarity is low, this order leads to high ratings. When object and verb are not similar, but verb and adverbial are, this order leads to low ratings. The strong effect of adverbial-verb similarity for -inchoatives when they follow the direct object and the smaller effect for the +inchoatives is in line with hypothesis A. The model is also in line with hypothesis C: when object-adverbial similarity is high, in the four top panels, -inchoatives are clearly associated with low ratings in pre-object position, but there is little or no effect or even improvement for +inchoatives, plausibly because of the availability of non-manner readings in this position.

### Experiment Ib

Experiment Ib employed a sentence fragment paradigm in which the participants could arrange the constituents in their preferred order (Gauza, Störzer, & Hörnig, 2020). The choice between **adverbial > object** and **object > adverbial** was modeled via logistic regression. Using the original predictors leads to a model including only **object status** as significant predictor: definites lead to a higher probability for the order **object > adverbial**, whereas indefinites don't lead to a significant preference. The factor **inchoative** does not emerge as significant. With the similarities, the effect of object status remains unchanged. All other factors, including **inchoative** participate in a fourway interaction ( $z = 2.628, p < 0.01$ ), plotted in Figure 3. Contra hypothesis A, we observe lit-

tle influence of the similarities on the ordering preferences for -inchoatives. For +inchoatives, when the similarity between direct object and verb is low (left-hand panels), the effect of the adverbial-verb similarity is reversed depending on the object-adverbial similarity: When it is high, adverbial-verb similarity positively correlates with the probability of the order **adverbial > object**. When it is low, the higher the adverbial-verb similarity, the greater the probability of the order **object > adverbial**. When the object-verb similarity is high, these effects are less distinct, but reversed in their general tendencies. As in Experiment Ia, object status and object-verb similarity do not interact, contra hypothesis B. Object-adverbial similarity is in line with hypothesis C in so far as it shows the clear interaction with inchoatives, however, not all effects are in line with the proposed competition account, and in fact, for +inchoatives, the model predicts the highest preference for **adverbial > object** both when all similarities are high, and when only the adverbial-verb similarity is high.

### Experiment II

Experiment II is similar to Ia except that each target sentence was preceded by a question about the direct object. The direct object was therefore given, aiming to control possible effects of information structure on ordering. Including the original predictors, only order itself emerges as significant: **object > adverbial** is correlated with higher naturalness ratings. In the model with similarities, predictor order remains significant, but the model also contains an interaction between **inchoative** and adverbial-verb similarity ( $t = -2.859, p = 0.01$ ). Adverbial-verb similarity does not correlate with the ratings

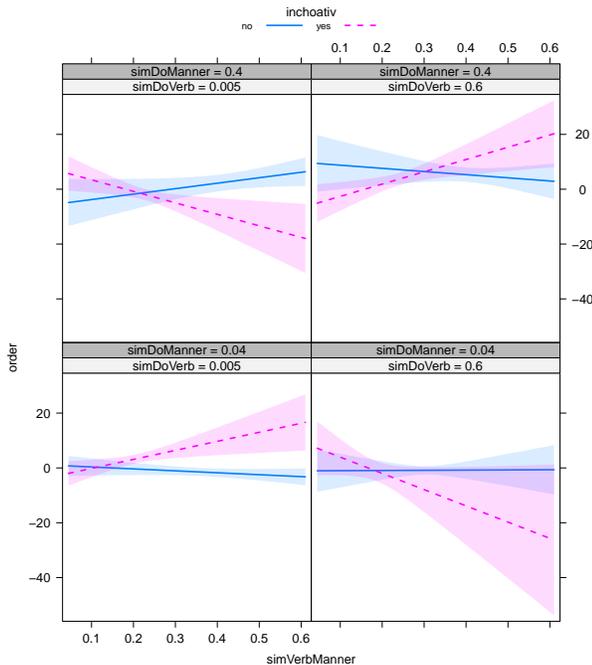


Figure 3: Experiment Ib, predictor effect plot for **inchoative** and all interactions.

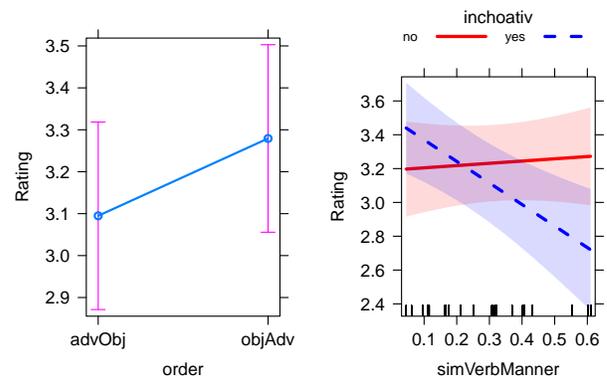


Figure 4: Experiment II, predictor effect plots for **order** and **inchoative** with their respective interactions.

for the -inchoatives, for +inchoatives, the higher the similarity, the lower the ratings. Again, this is not in line with hypothesis A. As **object status** did not become significant in the model with the original predictors, it is perhaps not surprising that the object-verb similarity also did not emerge as a significant predictor. Note that there is little additional variance explained by the model, in contrast to the naturalness ratings in Experiment Ia.

### Experiment III, IVa, and IVb

Experiment III was again a naturalness judgement task, target sentences varying whether [object + verb] or [adverbial + verb] were preposed. Without the distributional predictors, **order** (that is, preposing either adverbial or object) enters in two-way interactions with both **object status** and **inchoative**. With distributional predictors, both factors participate in further interactions. This is the only experiment in which **object status** directly interacts with distributional similarities (both object-verb and adverbial-verb similarity). Notably, in these interactions the clear across-the-board facilitation of +definites for sentences with preposed [adverbial + verb], and for -inchoatives in the same condition, observable in the non-distributional model does not obtain anymore.

Both Experiment IVa and IVb returned to the sentence structure from experiment Ia, with all direct objects realized as indefinite w-pronoun *was* ‘what’. Modeling the naturalness ratings of IVa with the original predictors yields an interaction between **order** and **inchoative**: -inchoatives show no effect, but +inchoatives are judged better when preceding the direct object and worse when following it. With the distributional predictors, the interaction between inchoative and order remains in the model, but the model itself is completely dominated by the interactions of the three similarity measures, both with each other as well as with **order** and **inchoative** (cf. also the huge difference between the marginal pseudo R-squared values of these to models, 2.1% vs. 9.5%). For IVb, the self-paced reading study, we find no significant predictors for the critical region of indefinite + object in either order. For the spill-over region containing the verb, we find a small but significant effect of object-verb similarity ( $t = -2.135$ ,  $p = 0.0406$ ): the higher the similarity, the faster the reading time.

### General discussion and conclusion

Across all six experiments, we arrived at significantly better models with broader coverage. The interaction of distributional predictors not only with the lexical predictor **inchoative** (except IVb) but also **order** (in Ia, Ib, III, and IVa) and **object status** (III) shows that these predictors are needed to understand the subjects’ ordering preferences. Already the first three experiments make this abundantly clear: Without the distributional predictors, the factor **inchoative** is only significant in Experiment Ia, not in Ib and II. With distributional predictors, we see that this factor plays a significant role across all three experiments, always modulated via the similarities. In addition, the effect of **object status** becomes much more nuanced when the similarities are taken into account, most striking in the results of Experiment Ia, where, for the -inchoatives, any difference between its two levels vanishes and the clear association of **object** > **adverbial** with higher ratings originally found turns out to be completely dependent on object-adverbial and verb-manner similarities, which modulate its effect across the whole spectrum of ratings. Taken together, the remodeling results show only weak

support for the assumption of a default position of manner adverbials per se, in contrast to the interpretation of the experimental results that suggests itself when only considering the original predictors.

These results show that using the cosine similarities between high-dimensional distributional vectors as stand-in for shared experience of given constituents and thus as a correlate for ease of conceptual integration is a very promising enterprise. However, the complex interactions emerging in most models show that the hypotheses brought forward at the beginning were too simplistic: neither did the inchoative consistently mute the effect of adverbial-verb similarity, nor did object status show much interaction with object-verb similarity. And while manner-object similarity emerged as a significant factor in several of the models, it again was not easily interpretable due to its participation in complex interactions. On the whole, the massive interaction observed is not surprising when considering the original motivation behind this reanalysis: if both objects and adjuncts can integrate, and if similarities are stand-ins for the likelihood of integration, we expect the resulting default position to be determined by the interactions of their respective probabilities, among themselves and with other factors. Perhaps more often than not, these interactions may result in situations in which these probabilities are so finely balanced that yet other factors influence the eventual decision for an order.

The most pressing challenge for future work is to further tease apart the role of object-adverbial similarity in terms of competition and complementation. One promising approach for this is the addition of purely collocation based measures such as PMI in order to capture the likelihood of cooccurrence and the relative ordering of these two constituents. Equally of interest in this respect are vector-composition based models. Both approaches have been successfully applied in the domain of attributive adjective ordering, cf. Hahn, Degen, Goodman, Jurafsky, and Futrell (2018) and Vecchi, Zamparelli, and Baroni (2013), respectively.

### Acknowledgments

Many thanks to Fabian Schlotterbeck for his insightful and very helpful comments on an earlier version of this paper. Work on this paper was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – SFB 833 – Project ID 75650358.

### References

- Bartoń, K. (2016). Mumin: Multi-model inference [Computer software manual]. Retrieved from <http://CRAN.R-project.org/package=Mumin> (R package version 1.15.6)
- Eckardt, R. (2003). Manner adverbs and information structure. Evidence from the adverbial modification of verbs of creation. In E. Lang, C. Fabricius-Hansen, & C. Maienborn (Eds.), *Modifying adjuncts* (pp. 261–305). Berlin: Mouton de Gruyter.

- Ernst, T. (2002). *The syntax of adjuncts*. Cambridge: Cambridge University Press.
- Fox, J. (2003). Effect displays in R for generalised linear models. *Journal of Statistical Software*, 8(15), 1–27.
- Frey, W., & Pittner, K. (1998). Zur Positionierung der Adverbiale im Deutschen Mittelfeld. *Linguistische Berichte*(176), 489–534.
- Gauza, H. (2018). *Position und Verarbeitung VP-interner Adjunkte*. Dissertation, Universität Tübingen. doi: 10.15496/publikation-24440
- Gauza, H., Störzer, M., & Hörnig, R. (2020). *The role of competition: An evaluation of sentence fragment arrangement as a method for determining linearization preferences*. (In revision)
- Günther, F., Rinaldi, L., & Marelli, M. (2019). Vector-space models of semantic representation from a cognitive perspective: A discussion of common misconceptions. *Perspectives on Psychological Science*, 14(6), 1006–1033.
- Hahn, M., Degen, J., Goodman, N., Jurafsky, D., & Futrell, R. (2018). An information-theoretic explanation of adjective ordering preferences. In *Proceedings of the 40th annual meeting of the cognitive science society* (pp. 1766–1771).
- Jacobs, J. (1993). Integration. In M. Reis (Ed.), *Wortstellung und Informationsstruktur* (pp. 63–116). Niemeyer.
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software*, 82(13), 1–26.
- Lebani, G. E., & Lenci, A. (2018). A distributional model of verb-specific semantic roles inferences. In T. Poibeau & A. Villavicencio (Eds.), *Language, cognition, and computational models* (p. 118–158). Cambridge University Press.
- McRae, K., Ferretti, T. R., & Amyote, L. (1997). Thematic roles as verb-specific concepts. *Language and Cognitive Processes*, 12(2-3), 137–176.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient Estimation of Word Representations in Vector Space. *ArXiv e-prints*.
- Schäfer, M. (2013). *Positions and interpretations: German adverbial adjectives at the syntax-semantics interface*. Berlin: De Gruyter Mouton.
- Vecchi, E. M., Zamparelli, R., & Baroni, M. (2013). Studying the recursive behaviour of adjectival modification with compositional distributional semantics. In *Proceedings of emnlp 2013* (pp. 141–151). Seattle: ACL.