

# Combining Abstractness and Language-specific Theoretical Indicators for Detecting Non-Literal Usage of Estonian Particle Verbs

Eleri Aedmaa<sup>1</sup>, Maximilian Köper<sup>2</sup> and Sabine Schulte im Walde<sup>2</sup>

<sup>1</sup> Institute of Estonian and General Linguistics, University of Tartu, Estonia

<sup>2</sup> Institut für Maschinelle Sprachverarbeitung, Universität Stuttgart, Germany

elери.aedmaa@ut.ee

{maximilian.koeper, schulte}@ims.uni-stuttgart.de

## Abstract

This paper presents two novel datasets and a random-forest classifier to automatically predict literal vs. non-literal language usage for a highly frequent type of multi-word expression in a low-resource language, i.e., Estonian. We demonstrate the value of language-specific indicators induced from theoretical linguistic research, which outperform a high majority baseline when combined with language-independent features of non-literal language (such as abstractness).

## 1 Introduction

Estonian particle verbs (PVs) are multi-word expressions combining an adverbial particle with a base verb (BV), cf. Erelt et al. (1993). They are challenging for automatic processing because their components do not always appear adjacent to each other, and the particles are homonymous with adpositions. In addition, as illustrated in examples (1a) vs. (1b), the same PV type can be used in literal vs. non-literal language.

- (1) a. Ta astu-s kaks sammu tagasi.  
he step-PST.3SG two step.PRT back  
'He took two steps back.'
- b. Ta astu-s ameti-st tagasi.  
he step-PST.3SG job-ELA back  
'He resigned from his job.'

Given that the automatic detection of non-literal expressions (including metaphors and idioms) is critical for many NLP tasks, the last decade has seen an increase in research on distinguishing literal vs. non-literal meaning (Birke and Sarkar, 2006, 2007; Sporleder and Li, 2009; Turney et al., 2011; Shutova et al., 2013; Tsvetkov et al., 2014; Köper and Schulte im Walde, 2016).

Most research up to date has, however, focused on resource-rich languages (mainly English and German), and elaborated on general indicators – such as contextual abstractness – to identify non-literal language. As to our knowledge, only Tsvetkov et al. (2014) and Köper and Schulte im Walde (2016) explored language-specific features.

The aim of this work is to automatically predict literal vs. non-literal language usage for a very frequent type of multi-word expression in a low-resource language, i.e., Estonian. The predicate is the center of grammatical and usually semantic structure of the sentence, and it determines the meaning and the form of its arguments, cf. Erelt et al. (1993). Hence, the surrounding words (i.e., the context), their meanings and grammatical forms could help to decide whether the PV should be classified as compositional or non-compositional.

In addition to applying language-independent features of non-literal language, we demonstrate the value of indicators induced from theoretical linguistic research, that have so far not been explored in the context of compositionality. For this purpose, this paper introduces two novel datasets and a random-forest classifier with standard and language-specific features.

The remainder of this paper is structured as follows. We give a brief overview of previous studies on Estonian PVs in Section 2, and Section 3 introduces the target dataset. All features are described in Section 4. Section 5 lays out the experiments and evaluation of the model, and we conclude our work in Section 6.

## 2 Related Work

The compositionality of Estonian PVs has been under discussion in the theoretical literature for

decades but still lacks a comprehensive study. [Tragel and Veismann \(2008\)](#) studied six verbal particles and their aspectual meanings, and described how horizontal and vertical dimensions are represented. [Veismann and Sahkai \(2016\)](#) investigated the prosody of Estonian PVs, finding PVs expressing perfectivity the most problematic to classify.

Recent computational studies on Estonian PVs involve their automatic acquisition ([Kaalep and Muischnek, 2002](#); [Uiboaed, 2010](#); [Aedmaa, 2014](#)), and predicting their degrees of compositionality ([Aedmaa, 2017](#)). [Muischnek et al. \(2013\)](#) investigated the role of Estonian PVs in computational syntax, focusing on Constraint Grammar. Most research on automatically detecting non-literal language has been done on English and German (as mentioned above), and elaborated on general indicators to identify non-literal language. Our work is the first attempt to automatically distinguish literal and non-literal usage of Estonian PVs, and to specify on theory- and language-specific features.

### 3 Target PV Dataset

For creating a dataset of literal and non-literal language usage for Estonian PVs, we selected 210 PVs across 34 particles: we started with a list of 1,676 PVs that occurred at least once in a 170-million token newspaper subcorpus of the Estonian Reference Corpus<sup>1</sup> (ERC) and removed PVs with a frequency  $\leq 9$ . Then we sorted the PVs according to their frequency and selected PVs across different frequency ranges for the dataset. In addition, we included the 20 most frequent PVs. We plan to analyse the influence of frequency on the compositionality of PVs in future work, thus it was necessary to collect evaluations for PVs with different frequencies.

For each of the 210 target PVs, we then automatically extracted 16 sentences from the ERC. The sentences were manually double-checked to make sure that verb and adverb formed a PV and did not appear as independent word units in a clause. The choice of the numbers of PVs and sentences relied on the fact of limited time and other resources that allowed us to evaluate approximately 200 PVs and 2,000 sentences.

<sup>1</sup>[www.cl.ut.ee/korpused/segakorpus/](http://www.cl.ut.ee/korpused/segakorpus/)

The resulting set of sentences was evaluated by three annotators with a linguistic background. They were asked to assess each sentence by answering the question: "What is the usage of the PV in the sentence on a 6-point scale ranging from clearly literal (0) to clearly non-literal (5) language usage?" In case of multiple PVs in the same sentence, the information of which PV to evaluate was provided for the annotators. Although we use binary division of PVs in this study, it was reasonable to collect evaluations on a larger than binary scale because of the following reasons: first, it is a well-known fact that multi-word expressions do not fall into the binary classes of compositional vs. non-compositional expressions ([Bannard et al., 2003](#)), and second, it was important to create a dataset that would be applicable to multiple tasks. Thus our dataset can be used to investigate the degrees of compositionality of PVs in the future.

The agreement among 3 annotators on all 6 categories is fair (Fleiss'  $\kappa = 0.36$ ). A binary distinction based on the average sentence scores into literal (average  $\leq 2.4$ ) and non-literal (average  $\geq 2.5$ ) resulted in substantial agreement ( $\kappa = 0.73$ ). Our experiments below use the binary-class setting, disregarding all cases of disagreement.

This final dataset<sup>2</sup> includes 1,490 sentences: 1,102 non-literal and 388 literal usages across 184 PVs with 120 different base verbs and 32 particle types. 63 PVs occur only in non-literal sentences, 15 only in literal sentences and 106 PVs in non-literal and literal sentences. From 120 verbs 50 appear only in non-literal sentences, 15 only in literal sentences, and 55 verbs in both literal and non-literal sentences. The distribution of (non-) literal sentences across particle types is shown in [Figure 1](#). While many particles appear mostly in non-literal language (and *esile*, *alt*, *ihte*, *ära* are exclusively used in their non-literal meanings in our dataset), they all have literal correspondences. No particle types appear only in literal sentences.

### 4 Features

In this section we introduce standard, language-independent features (unigrams and abstractness) as well as language-specific features (case and animacy) that we will use to distinguish literal and non-literal language usage of Estonian PVs.

<sup>2</sup>The dataset is accessible from <https://github.com/elieriaedmaa/compositionality>.




Figure 1: (Non-)literal language usage across particles.

**Unigrams** Our simplest language-independent features are unigrams, i.e., lemmas of content words that occur in the same sentences with our target PVs. More precisely, unigrams are the list of lemmas of all words that we induced from all our target sentences (there is at least one PV in each sentence), after excluding lemmas that occurred  $\leq 5$  times in total.

**Abstractness** Abstractness has previously been used in the automatic detection of non-literal language usage (Turney et al., 2011; Tsvetkov et al., 2014; Köper and Schulte im Walde, 2016), as abstract words tend to appear in non-literal sentences. Since there were no ratings for Estonian, we followed Köper and Schulte im Walde (2016) to automatically generate abstractness ratings for Estonian lemmas: we translated 24,915 English lemmas from Brysbaert et al. (2014) to Estonian relying on the English-Estonian Machine Translation dictionary<sup>3</sup>. We then lemmatized the 170-million token ERC subcorpus and created a vector space model. To learn word representations, we relied on the skip-gram model from Mikolov et al. (2013). Finally, we applied the algorithm from Turney et al. (2011) using the 29,915 translated ratings from Brysbaert et al. (2014) as seeds.

<sup>3</sup><http://www.eki.ee/dict/ies/>

This algorithm relies on the hypothesis that the degree of abstractness of a word’s context is predictive of whether the word is used in a metaphorical or literal sense. The algorithm learns to assign abstractness scores to every word representation in our vector space, resulting in a novel resource<sup>4</sup> of automatically created ratings for 243,675 Estonian lemmas.

Unfortunately we can not provide an evaluation for this dataset at the moment, because Estonian is lacking a suitable human-judgement-based gold standard. In addition, the creation would require extensive psycholinguistic research which falls far from the authors’ specialization.

We adopted the following abstractness features from Turney et al. (2011) and Köper and Schulte im Walde (2016): average rating of all words in a sentence, average rating of all nouns in a sentence (including proper names), rating of the PV subject, and rating of the PV object.

The ratings of PV subject and object express the abstractness score of the head of the noun phrase. For example, the average score of the object (i.e., *oma koera*) in the sentence (2c) is the rating of the head of the noun phrase (i.e., *koer*), not the average of the ratings of the determiner and the head.

<sup>4</sup>The dataset is accessible from <https://github.com/eleriaedmaa/compositionality>.

We assume that the subjects and objects are more concrete in literal sentences. For example, the subject (*sõber*) and the object (*koer*) in the literal sentences (2a) and (2c) are more concrete than the subject (*surm*) and object (*viha*) in the non-literal sentences (2b) and (2d).

- (2) a. **Sõber** jook*s-i-s* mu-*lle* järele.  
**friend** run-PST-3SG I-ALL after  
 ‘A friend ran after me.’
- b. **Surm** jook*s-i-s* ta-*lle* järele.  
**death** run-PST-3SG he-ALL after  
 ‘The death ran after him.’
- c. Mees suru-*s* **koera** maha.  
 man push-PST.3SG **dog.GEN** down  
 ‘The man pushed the dog down.’
- d. Mees suru-*s* **viha** maha.  
 man push-PST.3SG **anger.GEN** down  
 ‘The man suppressed his anger.’

Figure 2 illustrates the abstractness scores for literal vs. non-literal sentences. In general, literal sentences are clearly more concrete, especially when looking at nouns only, and even more so when looking at the nouns in specific subject and object functions.



Figure 2: Abstractness scores across literal and non-literal sentences.

**Subject and object case** Estonian distinguishes between “total” subjects in the nominative case and “partial” subjects in the partitive case. Partial subjects are not in subject-predicate agreement (Erelt et al., 1993). For example, the subject *külaline* receives nominative case in sentence (3b) and partitive case in sentence (3a). We observed that subject case assignment often correlates with (non-)literal readings; in the examples, sentence (3a) is literal, and sentence (3b) is non-literal.

- (3) a. **Külaline** tule-*b* juurde.  
**guest.NOM** come-3SG near  
 ‘The guest approaches.’
- b. **Külali-si** tule-*b* juurde.  
**guest.PL.PRT** come-3SG up  
 ‘The number of guests is increasing.’

Similarly, a “total” object in Estonian receives nominative or genitive case, and a “partial” object receives partitive case. For example, the object *supi* in sentence (4a) is assigned genitive case, and the object *mida* in sentence (4b) partitive case. In sentence (4a), the meaning of the PV *ette võtma* is literal; in sentence (4b) the meaning is non-literal.

- (4) a. Tüdruk võt-*tis* ette **supi**.  
 girl take-PST.3SG front **soup.GEN**  
 ‘The girl took the soup in front of her.’
- b. **Mida** koos ette võt-*ta?*  
**what.SG.PRT** together front take-INF  
 ‘What should we do together?’

Figure 3 illustrates that the distribution of subject and object cases across literal and non-literal sentences does not provide clear indicators. In addition, the correlation between subject/object case and (non-)literalness has not been examined thoroughly in theoretical linguistics. But based on corpus analyses as exemplified by the sentences (3b)–(4b), we hypothesize that the case distribution might provide useful indicators for (non-)literal language usage.

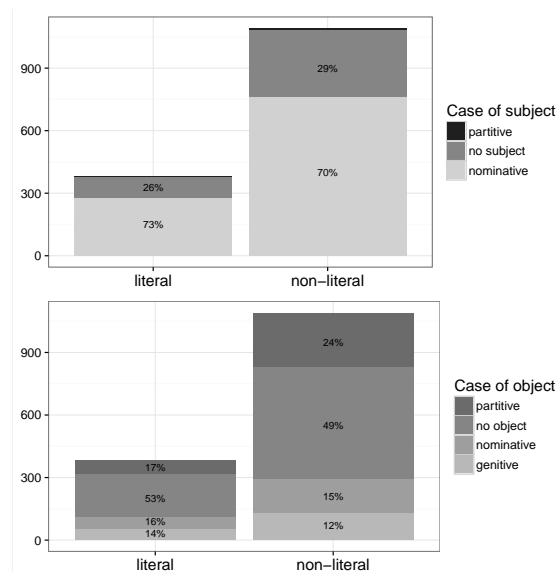


Figure 3: Distribution of subject/object case across literal and non-literal sentences.

**Subject and object animacy** According to Estonian Grammar (Erelt et al., 1993) the meaning of the predicate might determine (among other features) the animacy of its arguments in a sentence. If the verb requires an animate subject, but the subject is inanimate, the meaning of the sentence is non-literal. For example, the PV in the sentences (5a) and (5b) is the same (*sisse kutsuma* 'to invite in'), but in the first sentence the subject *sõber* is animate and the sentence is literal, while the subject *maja* in sentence (5b) is inanimate and the sentence is non-literal. Similarly, the subject *naine* in sentence (5c) is animate and the sentence is literal, while the subject *välimus* in sentence (5d) is inanimate and the sentence is non-literal.

As before, the correlation between subject animacy and (non-)literalness has not been examined thoroughly in theoretical linguistics, but the animacy of the subject seems to correlate with the (non-)literalness of the sentences.

- (5) a. **Sõber** kutsu-s mu sisse.  
**friend** invite-PST.3SG I.GEN in  
 'A friend invited me in.'
- b. **Maja** ei kutsu sisse.  
**house** NEG invite.CONNEG in  
 'The house doesn't look inviting.'
- c. **Naine** tõuka-s mehe eemale.  
**woman** push-PST.3SG man.GEN away  
 'A woman pushed a man away.'
- d. Poe **välimus** tõuka-s  
 shop.GEN **appearance** push-PST.3SG  
 mehe eemale.  
 man.GEN away  
 'The appearance of the shop made the man go away.'

The impact of the object animacy on the meaning of the PVs is less intuitive, but still the object in sentence (6a) is inanimate and the meaning of the PV is literal, while the object in sentence (6b) is animate and the meaning of the PV is non-literal.

- (6) a. Mees põleta-s **kaitsme** läbi.  
 man burn-PST.3SG fuse.GEN out  
 'The man burned the fuse.'
- b. Mees põleta-s **enese** läbi.  
 man burn-PST.3SG himself.GEN out  
 'The man had a burnout.'

There are no explicit connections between the subject animacy pointed out in the literature. Figure 4 shows the distribution of animacy across subjects and objects across the literal and non-literal usage. The differences in numbers are not remarkable, but based on the examples, we assume that the animacy of the subject might have an impact on the literal and non-literal usage of PVs. Thus, we include animacy into our feature space.




Figure 4: Distribution of subject/object animacy across literal and non-literal language.

Sentences (6a) and (6b) demonstrate that the abstractness/concreteness scores may already indicate the (non-)literal usage of the PV and the feature of animacy does not add any information: the concrete words are inanimate and they appear in the literal sentences, and the animate (and abstract) words in non-literal sentences. Still, as shown in sentences (5a)–(5d), the concrete subject of literal sentence can be also animate (i.e., *sõber*, *naine*), the concrete subject of non-literal sentence can be inanimate (i.e., *maja*), and the inanimate subject of non-literal sentence might be abstract (i.e., *välimus*). Thus, we argue that the abstractness ratings are not sufficient to express the animacy of the words and animacy can be useful as feature for the detection of (non-)literal usage of Estonian PV.

**Case government** Case government is a phenomenon where the lexical meaning of the base verb influences the grammatical form of the argument, e.g., the predicate determines the case of the

argument (Erelt et al., 1993). Thus, argument case depends on the meaning of the PV. For example, in sentence (7a) the PV *läbi minema* ‘to go through’<sup>5</sup> is literal and requires an argument that answers the question *from where?* Hence, the argument has to receive relative case. In sentence (7b) the PV provides a non-literal meaning (‘to succeed’) and does not require any additional arguments. We hypothesize that the case of the argument is helpful to predict (non-)literal usage of PVs.

- (7) a. Ta läks metsa-st läbi.  
 she go.PST.3SG forest-SG.EL through  
 ‘She went through the forest.’
- b. Mu ettepanek läks läbi.  
 I.GEN proposal go.PST.3SG through  
 ‘My proposal was successful.’

Note that the cases of the subject and object are individual features in our experiments, the feature of case government includes the cases of other types of arguments, i.e., adverbials and modifiers.

In addition, Figure 5 introduces the distribution of the argument case across the literal and non-literal sentences, and shows that not all cases (e.g., inessive, translative) appear in both types of sentences.




Figure 5: Distribution of argument case across literal and non-literal sentences.

Compared to all other features described in this section, animacy is the most problematic because the information is not obtained automatically. For

<sup>5</sup>The verb *minema* ‘to go’ is irregular and the stem is not derivable from the infinitive.

the abstractness scores we use the previously described dataset, and the cases of subjects, objects and other arguments are accessible with the help of the morphological analyser<sup>6</sup> and the part-of-speech tagger<sup>7</sup>. At the moment, the animacy information about the subject and object are added manually by the authors.

## 5 Experiments and Results

The classification experiments to distinguish between literal and non-literal language usage of Estonian PVs rely on the sentence features defined above. They were carried out using a random forest classifier (Breiman, 2001) that constructs a number of randomized decision trees during the training phase and makes prediction by averaging the results. For our experiments, we used 100 random decision trees. The random forest classifier performs better in comparison of other classification methods that we have applied in the Weka toolkit (Witten et al., 2016). For the evaluation we perform 10-fold cross validation, hence we use the previously described data for training and testing.

The classification results across features and combinations of features are presented in Table 1. We report accuracy as well as  $F_1$  for literal and non-literal sentences.

Table 1 shows that the best single feature types are the unigrams (acc: 82.3%) and the base verbs (81.2%). Combining the two, the accuracy reaches 84.2%. No other single feature type goes beyond the high majority baseline (74.0%), but the combinations in the Table 1 significantly outperform the baseline, according to  $\chi^2$  with  $p < 0.01$ .

Adding the particle type to the base verb information (1–2) correctly classifies 85.2% of the sentences. Further adding unigrams (1–3), however, does not help. Regarding abstractness, adding the ratings for all but objects to the particle-verb information (1–2, 4–6) is best and reaches an accuracy of 86.3%. Subject case information, animacy and case government in combination with 1–2 reach similar values (85.3–86.3%). The overall best result (87.9%) is reached when combining particle and base verb information with all-noun and subject abstractness ratings, subject case, subject animacy, and case government.

<sup>6</sup>[http://www.filosoft.ee/html\\_morf\\_et/](http://www.filosoft.ee/html_morf_et/)

<sup>7</sup><http://kodu.ut.ee/~kaili/parser/>

feature type	acc	$F_1$	
		n-lit	lit
majority baseline	74.0%	85.0	0.00
1 particle (p)	73.6%	84.4	13.6
2 base verb (v)	81.2%	87.9	58.0
3 unigrams, $f > 5$ (uni)	82.3%	89.0	54.6
4 average rating of words (abs)	68.1%	79.8	24.5
5 average rating of nouns (abs)	68.5%	79.7	30.1
6 rating of the PV subject (abs)	72.3%	83.1	23.7
7 rating of the PV object (abs)	73.0%	83.5	25.2
8 subject case (case)	74.0%	85.0	0.00
9 object case (case)	74.0%	85.0	0.00
10 subject animacy (animacy)	74.0%	85.0	0.00
11 object animacy (animacy)	74.0%	85.0	0.00
12 case government (govern)	73.8%	84.6	10.1
p+v, 1–2	85.2%	90.3	68.7
v+uni, 2–3	84.2%	89.6	66.4
p+v+uni, 1–3	85.0%	90.1	68.5
p+v+abs, 1–2, 4–6	86.3%	90.9	72.3
p+v+abs, 1–2, 4–7	86.0%	90.7	71.3
p+v+abs, 1–2, 5–6	86.0%	90.7	71.9
p+v+case, 1–2, 8	85.3%	90.4	68.9
p+v+case, 1–2, 8–9	84.6%	89.7	69.3
p+v+animacy, 1–2, 10–11	86.2%	90.8	72.3
p+v+govern, 1–2, 12	86.2%	90.9	71.6
p+v+abs+lang, 1–2, 4–6, 10–12	87.3%	91.6	73.8
p+v+abs+lang, 1–2, 4–12	87.5%	91.8	73.8
p+v+abs+lang, 1–2, 5–6, 8, 10, 12	<b>87.9%</b>	<b>92.0</b>	<b>75.0</b>

Table 1: Overview of classification results.

While Table 1 only lists a selection of all possible combinations of features to present the most interesting cases, it illustrates that the combination of language-independent features and language-specific features is able to outperform the high majority baseline. Although the difference between the best combination without language-specific features (86.3%) and the best combination with language-specific features (87.9%) is not statistically significant, the best-performing combination provides  $F_1=92.0$  for non-literal sentences and  $F_1=75.0$  for literal sentences.

## 6 Conclusion

This paper introduced a new dataset with 1,490 sentences of literal and non-literal language usage for Estonian particle verbs, a new dataset of abstractness ratings for >240,000 Estonian lemmas across word classes, and a random-forest classifier that distinguishes between literal and non-literal sentences with an accuracy of 87.9%.

The most salient feature selection confirms our theory-based hypotheses that subject case, subject animacy and case government play a role in non-

literal Estonian language usage. Combined with abstractness ratings as language-independent indicators of non-literal language as well as verb and particle information, the language-specific features significantly outperform a high majority baseline of 74.0%.

## Acknowledgements

This research was supported by the University of Tartu ASTRA Project PER ASPERA, financed by the European Regional Development Fund (Eleri Aedmaa), and by the Collaborative Research Center SFB 732, financed by the German Research Foundation DFG (Maximilian Köper, Sabine Schulte im Walde).

## References

- Eleri Aedmaa. 2014. Statistical methods for Estonian particle verb extraction from text corpus. In *Proceedings of the ESSLLI 2014 Workshop: Computational, Cognitive, and Linguistic Approaches to the Analysis of Complex Words and Collocations*. Tübingen, Germany, pages 17–22.
- Eleri Aedmaa. 2017. Exploring compositionality of Estonian particle verbs. In *Proceedings of the ESSLLI 2017 Student Session*. Toulouse, France, pages 197–208.
- Colin Bannard, Timothy Baldwin, and Alex Lascarides. 2003. A statistical approach to the semantics of verb-particles. In *Proceedings of the ACL Workshop on Multiword Expressions: Analysis, Acquisition and Treatment*. Sapporo, Japan, pages 65–72.
- Julia Birke and Anoop Sarkar. 2006. A clustering approach for the nearly unsupervised recognition of nonliteral language. In *Proceedings of the 11th Conference of the European Chapter of the ACL*. Trento, Italy, pages 329–336.
- Julia Birke and Anoop Sarkar. 2007. Active learning for the identification of nonliteral language. In *Proceedings of the Workshop on Computational Approaches to Figurative Language*. Rochester, NY, pages 21–28.
- Leo Breiman. 2001. Random forests. *Machine learning* 45(1):5–32.
- Marc Brysbaert, Amy Beth Warriner, and Victor Kuperman. 2014. Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods* 64:904–911.

- Tiiu Ereht, Ülle Viks, Mati Ereht, Reet Kasik, Helle Metslang, Henno Rajandi, Kristiina Ross, Henn Saari, Kaja Tael, and Silvi Vare. 1993. *Eesti keele grammatika II. Süntaks. Lisa: kiri [The Grammar of the Estonian Language II: Syntax]*. Tallinn: Eesti TA Keele ja Kirjanduse Instituut.
- Heiki-Jaan Kaalep and Kadri Muischnek. 2002. Using the text corpus to create a comprehensive list of phrasal verbs. In *Proceedings of the 3rd International Conference on Language Resources and Evaluation*. Las Palmas de Gran Canaria, Spain, pages 101–105.
- Maximilian Köper and Sabine Schulte im Walde. 2016. Automatically generated affective norms of abstractness, arousal, imageability and valence for 350 000 German lemmas. In *Proceedings of the 10th International Conference on Language Resources and Evaluation*. Portoroz, Slovenia, pages 2595–2598.
- Maximilian Köper and Sabine Schulte im Walde. 2016. Distinguishing literal and non-literal usage of German particle verbs. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. San Diego, California, USA, pages 353–362.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26*. Lake Tahoe, Nevada, USA, pages 3111–3119.
- Kadri Muischnek, Kaili Müürisep, and Tiina Puolakainen. 2013. Estonian particle verbs and their syntactic analysis. In *Proceedings of Human Language Technologies as a Challenge for Computer Science and Linguistics: 6th Language & Technology Conference*. Poznan, Poland, pages 7–11.
- Ekaterina Shutova, Simone Teufel, and Anna Korhonen. 2013. Statistical Metaphor Processing. *Computational Linguistics* 39(2):301–353.
- Caroline Sporleder and Linlin Li. 2009. Unsupervised recognition of literal and non-literal use of idiomatic expressions. In *Proceedings of the 12th Conference of the European Chapter of the Association for Computational Linguistics*. Athens, Greece, pages 754–762.
- Iloona Tragel and Ann Veismann. 2008. Kuidas horisontaalne ja vertikaalne liikumissuund eesti keeles aspektiks kehastuvad? [Embodiment of the horizontal and vertical dimensions in Estonian aspect]. *Keel ja Kirjandus* 7:515–530.
- Yulia Tsvetkov, Leonid Boytsov, Anatole Gershman, Eric Nyberg, and Chris Dyer. 2014. Metaphor detection with cross-lingual model transfer. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*. Baltimore, Maryland, pages 248–258.
- Peter Turney, Yair Neuman, Dan Assaf, and Yohai Cohen. 2011. Literal and metaphorical sense identification through concrete and abstract context. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Edinburgh, UK, pages 680–690.
- Kristel Uibo. 2010. Statistilised meetodid murdekorpuse ühendverbide tuvastamisel [Statistical methods for phrasal verb detection in Estonian dialects]. *Eesti Rakenduslingvistika Ühingu aastaraamat* 6:307–326.
- Ann Veismann and Heete Sahkai. 2016. Ühendverbidest läbi prosoodia prisma [Particle verbs and prosody]. *Eesti Rakenduslingvistika Ühingu aastaraamat* 12:269–285.
- Ian H. Witten, Eibe Frank, Mark A. Hall, and Christopher J. Pal. 2016. *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann.