

VERBCROCEAN: A Repository of Fine-Grained Semantic Verb Relations for Croatian

Ivan Sekulić and Jan Šnajder

University of Zagreb, Faculty of Electrical Engineering and Computing
Text Analysis and Knowledge Engineering Lab
Unska 3, 10000 Zagreb, Croatia
{ivan.sekulic, jan.snajder}@fer.hr

Abstract

In this paper we describe VERBCROCEAN, a broad-coverage repository of fine-grained semantic relations between Croatian verbs. Adopting the methodology of Chklovski and Pantel (2004) used for acquiring the English VerbOcean, we first acquire semantically related verb pairs from a web corpus hrWaC by relying on distributional similarity of subject-verb-object paths in the dependency trees. We then classify the semantic relations between each pair of verbs as similarity, intensity, antonymy, or happens-before, using a number of manually-constructed lexico-syntactic patterns. We evaluate the quality of the resulting resource on a manually annotated sample of 1000 semantic verb relations. The evaluation revealed that the predictions are most accurate for the similarity relation, and least accurate for the intensity relation. We make available two variants of VERBCROCEAN: a coverage-oriented version, containing about 36k verb pairs at a precision of 41%, and a precision-oriented version containing about 5k verb pairs, at a precision of 56%.

Keywords: verb relations, semantic relations, relation acquisition, Croatian language

1. Introduction

Lexico-semantic resources such as WordNet (Miller, 1995) play an important role in many natural language processing (NLP) tasks. In particular, for applications that involve processing of events and states, resources that model verb semantics are of great importance. A number of such resources have been developed, most notably verb lexica FrameNet (Baker et al., 1998) and VerbNet (Kipper-Schuler, 2005), which model the predicate-argument relations. More recently, the interest in textual entailment (Dagan et al., 2013) has motivated the construction of large-scale verb entailment resources (Pekar, 2006; Szpektor and Dagan, 2008; Hashimoto et al., 2009).

VerbOcean (Chklovski and Pantel, 2004) is a broad-coverage repository of fine-grained semantic relations between English verbs. VerbOcean models five semantic relations between verbs: similarity, strength, antonymy, enablement, and happens-before. The resource has been acquired semi-automatically from the web using lexico-syntactic patterns, and covers about 3,500 verbs and 30,000 semantic relations. VerbOcean has been used for many NLP tasks, including event extraction (Mani et al., 2006), paraphrase detection, entailment recognition (Dagan et al., 2010; Mehdad et al., 2011) and contradiction detection (De Marneffe et al., 2008).

In this paper we apply the VerbOcean methodology to extract a broad-coverage repository of semantic relations between Croatian verbs. We essentially adopt the procedure of Chklovski and Pantel (2004) and evaluate the quality of the resulting resource on a manually annotated sample of semantic verb relations. Based on our insights, we create VERBCROCEAN, a freely available repository of fine-grained semantic verb relations for Croatian. We make available two versions of the resource: a coverage-oriented version containing about 36k relations and a precision-oriented version containing about 5k relations.

2. Related Work

Only a handful of semantic verb resources for Croatian are publicly available. The Croatian WordNet (CroWN) (Raffaelli et al., 2008) contains, alongside nouns and adjectives, 2318 verbs arranged into hypernym hierarchies. CroDeriV (Šojat et al., 2013) is a comprehensive morphological lexicon containing almost 14,000 Croatian verbs. In a follow-up work, Šojat and Srebačić (2014) studied the morphosemantic relations between Croatian verbs in CroWN. The only publicly available predicate semantics resource is the Croatian Verb Valence Lexicon (Preradović et al., 2009), which lists 1,739 verbs with 5,118 valence frames, organized in 173 syntactic-semantic classes. Less verb-specific, Šnajder et al. (2013) present a distributional memory for Croatian, which includes syntax-based distributional vectors for about 2M lemmas, including the verbs.

3. Verb Relations

As noted by Schulte im Walde and Melinger (2005), different resources incorporate different semantic verb relations, and there seems to be no exhaustive study on the range of possible verb relations. Princeton WordNet identifies synonymy, antonymy, troponymy, hypernymy, entailment, and cause. On the other hand, VerbOcean identifies similarity, strength, antonymy, enabled, and happens-before. Following (Chklovski and Pantel, 2004), we chose to extract the following four semantic verb relations.

Similarity. Similarity includes synonymous and near-synonymous verbs (e.g., *get* :: *acquire*), but also verbs that differ in the degree of action, intensity, or the manner in which the action is carried out. Examples of similar verbs extracted in VERBCROCEAN include *informirati* :: *educirati* (to inform :: to educate), *snimiti* :: *fotografirati* (to record :: to photograph), and *poboljšati* :: *unaprijediti* (to ameliorate :: to advance).

Intensity. Even though two verbs are similar, one can refer to a more intense and stronger action. A pair of verbs belongs to this relation if the first verb refers to the weaker than (but similar to) the action to which the second verb refers. Example relations from VERBCROCEAN are *raniti :: ubiti* (to wound :: to kill), *komentirati :: kritizirati* (to comment :: to criticize), *misliti :: znati*, (to think :: to know), and *udvostručiti :: utrostručiti* (to double :: to triple). Note that intensity implies similarity.

Antonymy. Antonymy holds between two verbs that express opposite meaning. Examples from VERBCROCEAN are *potvrditi :: demantirati* (to confirm :: to confute), *kupiti :: prodati* (to buy :: to sell), and *povećati :: smanjiti* (increase :: decrease).

Happens-before. Verbs in this relation refer to two temporally disjoint intervals or instances. The first verb describes an event that generally happens before the one described by the second verb. Examples of this relation in VERBCROCEAN are *napasti :: udariti* (to attack :: to hit) and *dijagnosticirati :: liječiti* (to diagnose :: to treat).

A pair of verbs can be in multiple semantic relations. However, some relations are mutually exclusive (e.g., antonymy and similarity). As intensity implies similarity, antonyms cannot stand in the intensity relation. The happens-before relation can appear alongside any other relation. For instance, one action (e.g. *to wound*) can happen before an action that is similar to it (*to torture*), more intense than it (*to kill*), or its antonym (*to heal*).

4. VerbCROcean

Following Chklovski and Pantel (2004), we acquire VerbCROcean in two steps: (1) the acquisition of semantically related verb pairs based on statistical co-occurrence and (2) verb pair relation classification of using lexico-syntactic patterns. Prior to this, we preprocess the corpus and extract the verbs. We next describe these steps in detail.

4.1. Corpus and Preprocessing

We extract VERBCROCEAN from the Croatian web-corpus hrWaC (Ljubešić and Erjavec, 2011), totaling 1.2 billion words. The corpus has been lemmatized and tagged using the tools of Agić et al. (2013), and parsed using the tools of Agić and Merkler (2013).

It is a well-known fact that web corpora suffer from a considerable amount of noise (e.g., misspellings, ungrammatical language, jargon). In addition to this, in our case lemmatization introduces a substantial amount of systematic errors. To not compromise the precision of VERBCROCEAN, we decided to compile a whitelist of verbs consisting of those verbs occurring at least 500 times in hrWaC. Additionally, we discarded all verbs not ending in the infinitive suffix *-ti* or *-ći*, as well as modal verbs *biti* (to be), *htjeti* (will), *morati* (must), *moći* (can), and *trebati* (shall). The final list contains 7,050 verbs, which is almost twice as much as the number of verbs in VerbOcean. We also compiled a whitelist of nouns as those that occur more than 100 times in the corpora, resulting in 80,218 nouns. The nouns are used in the acquisition of verb pairs, as described below.

4.2. Acquisition of Verb Pairs

Following Chklovski and Pantel (2004), we acquire the semantically related verb pairs using the DIRT (Discovery of Inference Rules from Text) algorithm of Lin and Pantel (2001). DIRT utilizes the distributional hypothesis (Harris, 1954), which posits that the similarity of words correlates with the similarity of their contexts, with the difference that the similarities are computed between paths in dependency trees rather than between words.

DIRT was originally used for extracting paraphrases, but as we are interested only in verbs, we apply DIRT to paths of the form $subjectX \leftarrow verb \rightarrow objectY$. Subjects and objects fill slots X and Y of a single path. To reduce the number of spurious paths, we allow as slot fillers only the nouns from our whitelist of nouns.

Essentially, two paths are considered similar if they tend to link the same set of words, i.e., if they share a large number of words in the corresponding slots. Intuitively, not all words are equally important. For example, $X=person$ is less informative than $X=lawyer$. Therefore, similarity between two paths is computed based on the similarity of their slot fillers, weighted using pointwise mutual information. Following Lin and Pantel (2001), let (p, s, w) denote that slot s (where $s=subjectX$ or $s=objectY$) of path p is filled by word w . Mutual information between a path slot and its filler is computed as:

$$mi(p, s, w) = \ln \left(\frac{|p, s, w| \times |*, s, *|}{|p, s, *| \times |*, s, w|} \right)$$

where $|p, s, w|$ denotes the frequency of a triple (p, s, w) in the corpus. Let $T(p, s)$ be a set of all words that fill the slot s of path p . The similarity of two same-type slots s in paths p_1 and p_2 is defined as:

$$ssim(p_1, p_2, s) = \frac{\sum_{w \in T(p_1, s) \cap T(p_2, s)} mi(p_1, s, w) + mi(p_2, s, w)}{\sum_{w \in T(p_1, s)} mi(p_1, s, w) + \sum_{w \in T(p_2, s)} mi(p_2, s, w)}$$

Finally, the similarity between two paths p_1 and p_2 is computed as the geometric mean of similarities of their subject and object slots, $ssim(p_1, p_2, subjectX)$ and $ssim(p_1, p_2, objectY)$, respectively.

For reasons of efficiency, we do not compute the similarity between all path pairs. Instead, for each path, we first discard the paths that share with it less than 1% of words in the same-type slots. We then compute its similarity with the rest of the paths and keep up to 50 most similar paths. As a result, we extracted a total of 5,963 unique verbs that are semantically related to another verb in our dataset. On average, each verb is related to 39 other verbs, yielding a total of 183,662 unique verb pairs with a potential semantic relation. As an example, Table 1 lists top 20 verbs extracted as potentially related to the verb *voziti* (to drive).

4.3. Relation Classification

After acquiring the related verb pairs using DIRT, we classify the semantic relation of each pair by applying lexico-syntactic patterns on sentences from the corpus in which those pairs occur.

Similarity (3)	Antonymy (6)	Happens-before (15)	Intensity (5)
X tj, Y (<i>X i.e., Y</i>)	X ili Y (<i>X or Y</i>)	X * i onda * Y (<i>X * and then * Y</i>)	ne samo * X * već i * Y (<i>not only * X * but * Y</i>)
X i Y (<i>X and Y</i>)	X * tj, ne * Y (<i>X * i.e., not * Y</i>)	X * zatim * Y (<i>X * then * Y</i>)	X *, čak i * Y (<i>X *, and even * Y</i>)
X te Y (<i>X and Y</i>)	ne X i Y (<i>not X and Y</i>)	X pa Y (<i>X then Y</i>)	Y * ili barem * X (<i>Y * or at least * X</i>)
	X i ne Y (<i>X and not Y</i>)	X * i na kraju * Y (<i>X * and eventually * Y</i>)	Y * ili bar * X (<i>Y * or at least * X</i>)
	ne X te Y (<i>not X and Y</i>)	X * i poslije * Y (<i>X * and later * Y</i>)	X * a čak * Y (<i>X * and even * Y</i>)
	X te ne Y (<i>X and not Y</i>)	X * i kasnije * Y (<i>X * and later * Y</i>)	
		Y * nakon što * X (<i>Y * after * X</i>)	
		X * te onda * Y (<i>X * and then * Y</i>)	
		X * te na kraju * Y (<i>X * and eventually * Y</i>)	
		X * te kasnije * Y (<i>X * and later * Y</i>)	
		X * te poslije * Y (<i>X * and later * Y</i>)	
		prije * Y * i sada X (<i>before * Y * and now X</i>)	
		prije * Y * te sada X (<i>before * Y * and now X</i>)	
		X * potom * Y (<i>X * then * Y</i>)	
		Y, prije toga * X (<i>Y, before that * X</i>)	

Table 2: The 29 semantic relations and patterns that indicate their presence.

juriti (<i>to speed</i>)	parkirati (<i>to park</i>)
istrčati (<i>to run out</i>)	izletjeti (<i>to skid</i>)
prepriječiti (<i>to obstruct</i>)	naletjeti (<i>to hit</i>)
upaliti (<i>to turn on</i>)	sudariti (<i>to crash</i>)
upravljati (<i>to steer</i>)	sjesti (<i>to sit</i>)
odvoziti (<i>to drive off</i>)	letjeti (<i>to fly</i>)
približavati (<i>to approach</i>)	prevrnuti (<i>to roll over</i>)
dovezati (<i>to drive to</i>)	odvezati (<i>to untie</i>)
prevoziti (<i>to transport</i>)	zaustavljati (<i>to stop</i>)
skrenuti (<i>to turn</i>)	opkoliti (<i>to surround</i>)

Table 1: DIRT output for verb *voziti* (*to drive*).

Lexico-syntactic patterns. Table 2 lists the patterns for each of the four considered relations. We obtained 21 patterns by translating VerbOcean’s patterns to Croatian. However, as most English patterns do not translate well into Croatian, we decided to extend our list of patterns. To this end, we selected from the output of DIRT a list of 15 verb pairs that stand in one of the four considered relations. We then sampled sentences from hrWaC in which both verbs co-occur, and manually examined the syntactic paths between them. If the path seemed general enough, we add them to the list of patterns. In this way we were able to extract 8 more patterns from the corpus. We complemented the corpus search with web search using Google, but this did not result in any new patterns.

Relation classification. Given a verb pair extracted by DIRT, the classification of semantic relation is done as follows. For each sentence in the corpus that contains the given verb pair, we apply all of the above-listed 29 patterns as regular expressions, and then collect the frequencies of the verbs and the patterns. Based on these, we compute the strength of each relation using (1) the frequency of the relation, (2) pointwise mutual information (PMI), and (3) local mutual information (LMI). PMI is defined as follows:

$$PMI(v_1, rel, v_2) = \ln \left(\frac{P(v_1, rel, v_2)}{P(v_1)P(rel)P(v_2)} \right) \quad (1)$$

PMI is known to overestimate the strength of the association for rare words. To compensate for this, we also consider the frequency-weighted PMI, or local mutual information, proposed by Evert (2008):

$$LMI(v_1, rel, v_2) = f(v_1, rel, v_2) PMI(v_1, rel, v_2) \quad (2)$$

Relation refinement. Our model quite often predicts several relations for a given verb pair, each with a different confidence score. To make the predictions consistent with the definitions of semantic relations from Section 3, we post-process the relation predictions as follows. For each acquired relation, we take the confidence score (frequency, PMI, or LMI) of the verb pair (v_1, v_2) . We then refine the relations by applying the following four rules:

1. For symmetric relations (antonymy and similarity), we chose the one with the highest score, thereby considering also the score of the reverse pair (v_2, v_1) ;
2. If relation is classified as similarity, we also predict intensity, provided intensity was among the originally predicted relations;
3. If the relation is classified as intensity, but not as similarity, we also predict similarity with confidence equal to that of the intensity relation;
4. If the relation is happens-before, leave it as it is, because it can hold alongside any other relation.

5. Evaluation

In this section we describe a sample-based evaluation of VERBCROCEAN.

5.1. Verb Pairs Sample

We evaluate VERBCROCEAN on a sample of 1000 verb pairs manually annotated with semantic relations. We obtain the sample from the output of the relation classification step (but prior to the relation refinement step). We asked three annotators to independently annotate the sample by tagging each verb pair with one or more of the four semantic relations, or none if they think the verbs are not semantically related.

The inter-annotator agreement, measured in terms of Cohen’s kappa (Cohen, 1960), is 0.47, which is considered a moderate level of agreement according to Landis and Koch (1977). Note that we framed the problem as a multi-label task, and considered two annotations to be identical if and only if their sets of assigned labels are identical. As a consequence, the inter-annotator agreement estimate is probably somewhat pessimistic.

v_1	v_2	rel
izgubiti (<i>to lose</i>)	propustiti (<i>to miss</i>)	S
unijeti (<i>to bring in</i>)	osjećati (<i>to feel</i>)	N
požaliti (<i>to regret</i>)	obavijestiti (<i>to inform</i>)	N
posuditi (<i>to borrow</i>)	zaraditi (<i>to earn</i>)	A
govoriti (<i>to speak</i>)	komentirati (<i>to comment</i>)	S
plaćati (<i>to pay</i>)	nemati (<i>not to have</i>)	A
prići (<i>to approach</i>)	zaprijetiti (<i>to warn</i>)	B
naložiti (<i>to enjoin</i>)	zatražiti (<i>to require</i>)	S
regulirati (<i>to regulate</i>)	stabilizirati (<i>to stabilize</i>)	S
posjetiti (<i>to visit</i>)	kontaktirati (<i>to contact</i>)	S
umrijeti (<i>to die</i>)	dolaziti (<i>to come</i>)	N
pomoći (<i>to help</i>)	savjetovati (<i>to advise</i>)	S
spavati (<i>to sleep</i>)	sanjati (<i>to dream</i>)	BS
nadigrati (<i>to outplay</i>)	nadjačati (<i>to outperform</i>)	S
uzimati (<i>to take</i>)	uzeti (<i>to take away</i>)	BS
odustati (<i>to give up</i>)	pokušati (<i>to try</i>)	A
pobijediti (<i>to win</i>)	svladati (<i>to vanquish</i>)	S
širiti (<i>to spread</i>)	ograničavati (<i>to limit</i>)	A
udarati (<i>to hit</i>)	razbiti (<i>to break</i>)	BIS

Table 3: Excerpt from gold sample (*A* – *antonymy*, *S* – *similarity*, *B* – *hapens-before*, *I* – *intensity*, *N* – *no-relation*).

For the gold sample, we retain only the verb pairs for which at least two annotators agreed on their relation. We acquired a total of 953 verb pairs, distributed across the relations as follows: 73 antonymy, 151 happens-before, 15 intensity, 494 no-relation, and 261 similarity relations. An excerpt from the gold sample is presented in Table 3. We make the annotated sample freely available.¹

5.2. Results

The acquisition procedure has to trade-off between precision and coverage. To acquire a resource of satisfactory precision, but at the same time with the coverage as broad as possible, we set a threshold on each relation confidence score and acquire only the verb pairs for which the confidence score is above that threshold.

Fig. 1 shows the micro-averaged precision of the acquired semantic relations (measured on the manually annotated sample) with respect to the number of acquired verb pairs (determined by the confidence score threshold). Somewhat surprisingly, scoring the relations by frequency yields higher precision than scoring by PMI or LMI; we leave the investigation of this for future work. We also observe a general trend of precision decrease across the first 5,000 ranks, after which the precision stabilizes at about 0.4.

Based on these insights, we use frequency-based ranking to compile two variants of VERBCROCEAN: (1) a precision-oriented version (VCO-Pre), containing 5,010 verb pairs, with semantic relations classified at 56% micro-precision, and (2) a coverage-oriented version (VCO-Cov), containing all acquired relations (36,778 relations), with semantic relations classified at 41% micro-precision. We make both resources publicly available.¹

In Table 4 we show per class and macro accuracies for both VERBCROCEAN variants, measured in terms of precision, recall, and the F1-score on our manually annotated sample.




Figure 1: Precision of the semantic relations w.r.t. the number of acquired verb pairs.

	VCO-Pre			VCO-Cov		
	P	R	F1	P	R	F1
Similarity	0.57	0.20	0.30	0.56	0.63	0.59
Antonymy	0.62	0.11	0.19	0.47	0.23	0.31
Intensity	0	0	0	0.04	0.07	0.05
Happens-before	0.41	0.08	0.13	0.40	0.58	0.47
Macro	0.40	0.10	0.15	0.37	0.38	0.36

Table 4: Per class classification performance.

The predictions are most accurate for the similarity relation, and least accurate for the intensity relation. This is probably due to the fact that similarity is a more frequent relation, so many of its instances can be captured with the three patterns we used. At the same time, it is a rather broad relation, hence it can be acquired with a good precision (0.56). In contrast, intensity occurs rarely, and is also a more strict relation. Another relation that can be predicted with relatively good precision is antonymy. Overall, VCO-Pre has a macro-precision of 0.40 and a rather low macro-recall of 0.10. In contrast, VCO-Cov has well-balanced macro-precision and macro-recall of around 0.37. Note that, while we do report recall in Table 4, we do not measure the overall recall. Measuring the overall recall would require the annotation of a random sample of verb pairs from corpus. As most of the verbs stand in no semantic relation, annotating a sample of a satisfactory size would be too demanding.

6. Conclusion

We described the acquisition of VERBCROCEAN, a broad-coverage repository of semantic relations between Croatian verbs. We considered four semantic relations: synonymy, antonymy, intensity, and happens-before. Evaluation on a manually annotated sample reveals that we can extract 5,000 verb pairs with 56% micro-precision of semantic relation classification, and over 36K verb pairs with a micro-precision of 41%.

¹<http://takelab.fer.hr/verbcrocean>

v_1	v_2	rel	freq	pmi	lmi
pisati (to write)	čitati (to read)	S	744	15.15	48
poštovati (to respect)	voljeti (to love)	S	452	15.35	6936
demanirati (to deny)	potvrditi (to confirm)	A	398	17.05	6784
piti (to drink)	jesti (to eat)	S	398	17.14	6822
odbaciti (to reject)	prihvatiti (to accept)	A	244	16.59	4047
voljeti (to love)	mrziti (to hate)	A	174	17.57	3057
planirati (to plan)	organizirati (to organize)	S	147	13.39	1969
voziti (to drive)	udariti (to crash)	B	116	14.61	1695
dočekati (to welcome)	ugostiti (to host)	S	104	15.65	1628
diplomirati (to graduate)	doktorirati (to get a PhD)	B	100	18.41	1841
ispitati (to question)	privesti (to detain)	S	84	16.72	1404
pretući (to beat up)	silovati (to rape)	S	80	18.37	1469
blokirati (to block)	prosjedovati (to protest)	S	68	16.31	1109
prisustvovati (to attend)	sudjelovati (to partake)	S	64	12.15	849
povisiti (to raise)	sniziti (to lower)	A	60	19.55	1173
pobijediti (to win)	svladati (to overcome)	B	52	13.54	704
prihvatiti (to accept)	proslijediti (to forward)	B	18	13.77	247
zabiti (to score)	pobijediti (to win)	B	16	11.50	184
najaviti (to announce)	pokrenuti (to start up)	B	16	11.03	176
definirati (to define)	izgraditi (to construct)	B	14	13.14	184

Table 5: An excerpt from VERBCROCEAN.

There are a number of ways in which both precision and recall of the resource could be improved. Recall could be improved by including more patterns, which could be harvested in an automatic or semi-automatic manner. Precision could be improved by training a supervised model to give confidence-rated predictions for each verb pair, using a number of statistical and linguistic features. We intend to pursue this in future work.

7. Acknowledgments

This work has been supported by the Croatian Science Foundation under the project UIP-2014-09-7312.

8. Bibliographical References

- Agić, Ž. and Merkler, D. (2013). Three syntactic formalisms for data-driven dependency parsing of Croatian. In *Text, Speech, and Dialogue*, pages 560–567. Springer.
- Agić, Ž., Ljubešić, N., and Merkler, D. (2013). Lemmatization and morphosyntactic tagging of Croatian and Serbian. In *Proceedings of the 4th Biennial International Workshop on Balto-Slavic Natural Language Processing*, pages 48–57.
- Baker, C. F., Fillmore, C. J., and Lowe, J. B. (1998). The Berkeley FrameNet project. In *Proceedings of the 17th International Conference on Computational Linguistics-Volume 1*, pages 86–90.
- Chklovski, T. and Pantel, P. (2004). VerbOcean: Mining the web for fine-grained semantic verb relations. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 33–40.
- Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.
- Dagan, I., Dolan, B., Magnini, B., and Roth, D. (2010). Recognizing textual entailment: Rational, evaluation and approaches. *Natural Language Engineering*, 16(01):105–105.
- Dagan, I., Roth, D., Sammons, M., and Zanzotto, F. M. (2013). Recognizing textual entailment: Models and applications. *Synthesis Lectures on Human Language Technologies*, 6(4):1–220.
- De Marneffe, M.-C., Rafferty, A. N., and Manning, C. D. (2008). Finding contradictions in text. In *ACL*, volume 8, pages 1039–1047.
- Evert, S. (2008). Corpora and collocations. *Corpus Linguistics. An International Handbook*, 2.
- Harris, Z. S. (1954). Distributional structure. *Word*.
- Hashimoto, C., Torisawa, K., Kuroda, K., De Saeger, S., Murata, M., and Kazama, J. (2009). Large-scale verb entailment acquisition from the web. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 3-Volume 3*, pages 1172–1181.
- Kipper-Schuler, K. (2005). *VerbNet: A broad-coverage, comprehensive verb lexicon*. Ph.D. thesis. Computer and Information Science Dept., University of Pennsylvania.
- Landis, J. R. and Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, pages 159–174.
- Lin, D. and Pantel, P. (2001). Discovery of inference rules for question-answering. *Natural Language Engineering*, 7(04):343–360.
- Ljubešić, N. and Erjavec, T. (2011). hrWaC and slWaC: Compiling web corpora for Croatian and Slovene. In *Text, Speech and Dialogue*, pages 395–402. Springer.
- Mani, I., Verhagen, M., Wellner, B., Lee, C. M., and Pustejovsky, J. (2006). Machine learning of temporal relations. In *Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics*, pages 753–760. Association for Computational Linguistics.

- Mehdad, Y., Negri, M., and Federico, M. (2011). Using bilingual parallel corpora for cross-lingual textual entailment. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 1336–1345. Association for Computational Linguistics.
- Miller, G. A. (1995). WordNet: a lexical database for English. *Communications of the ACM*, 38(11):39–41.
- Pekar, V. (2006). Acquisition of verb entailment from text. In *Proceedings of the main conference on Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics*, pages 49–56.
- Preradović, N. M., Boras, D., and Kišiček, S. (2009). CROVALLEX: Croatian verb valence lexicon. In *Information Technology Interfaces, 2009. ITI'09. Proceedings of the ITI 2009 31st International Conference on*, pages 533–538.
- Raffaelli, I., Tadić, M., Bekavac, B., and Agić, Ž. (2008). Building Croatian WordNet. In *Proceedings of the 4th Global WordNet Conference*, pages 349–360.
- Schulte im Walde, S. and Melinger, A. (2005). Identifying semantic relations and functional properties of human verb associations. In *Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing*, pages 612–619. Association for Computational Linguistics.
- Šnajder, J., Padó, S., and Agić, Ž. (2013). Building and evaluating a distributional memory for Croatian. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, pages 784–789.
- Šojat, K. and Srebačić, M. (2014). Morphosemantic relations between verbs in Croatian WordNet. In *Proceedings of the Seventh Global WordNet Conference*, pages 262–267.
- Šojat, K., Srebačić, M., and Štefanec, V. (2013). CroDeriV and the morphological analysis of Croatian verbs. *Suvremena lingvistika*, 39(75):75–96.
- Szpektor, I. and Dagan, I. (2008). Learning entailment rules for unary templates. In *Proceedings of the 22nd International Conference on Computational Linguistics-Volume 1*, pages 849–856.