Likelihood of external causation in the structure of events

Tanja Samardžić
CorpusLab, URPP Language and Space
University of Zurich
`tanja.samardzic@uzh.ch`

Paola Merlo
Linguistics Department
University of Geneva
`Paola.Merlo@unige.ch`

Abstract

This article addresses the causal structure of events described by verbs: whether an event happens spontaneously or it is caused by an external causer. We automatically estimate the likelihood of external causation of events based on the distribution of causative and anticausative uses of verbs in the causative alternation. We train a Bayesian model and test it on a monolingual and on a bilingual input. The performance is evaluated against an independent scale of likelihood of external causation based on typological data. The accuracy of a two-way classification is 85% in both monolingual and bilingual setting. On the task of a three-way classification, the score is 61% in the monolingual setting and 69% in the bilingual setting.

1 Introduction

Ubiquitously present in human thinking, causality is encoded in language in various ways. Computational approaches to causality are mostly concerned with automatic extraction of causal schemata (Michotte, 1963; Tversky and Kahneman, 1982; Gilovich et al., 1985) from spontaneously produced texts based on linguistic encoding. A key to success in this endeavour is understanding how human language encodes causality.

Linguistic expressions of causality, such as causative conjunctions, verbs, morphemes, and constructions, are highly ambiguous, encoding not only the real-world causality, but also the structure of discourse, as well as speakers' attitudes (Moeschler, 2011; Zufferey, 2012). Causality judgements are hard to elicit in an annotation project. This results in a low inter-annotator agreement and makes the evaluation of automatic systems difficult (Bethard, 2007; Grivaz, 2012).

Our study addresses the relationship between world-knowledge about causality and the grammar of language, focusing on the causal structure of events expressed by verbs. In current analyses, the meaning of verbs is decomposed into multiple predicates which can be in a temporal and causal relation (Pustejovsky, 1995; Talmy, 2000; Levin and Rappaport Hovav, 2005; Ramchand, 2008).

(1) a. *Causative*: Adam broke the laptop.
     b. *Anticausative*: The laptop broke.

We propose a computational approach to the *causative alternation*, illustrated in (1), in which an event (*breaking the laptop* in (1)) can be dissociated from its immediate causer (*Adam* in (1a)). The causative alternation has been attested in almost all languages (Schafer, 2009), but it is realised with considerable cross-linguistic variation in the sets of alternating verbs and in the grammatical encoding (Alexiadou et al., 2006; Alexiadou, 2010).

Since the causative alternation involves most verbs, identifying the properties of verbs which allow them to alternate is important for developing representations of the meaning of verbs in general. Analysing the structural components of the meaning of verbs proves important for tasks such as word sense disambiguation (Lapata and Brew, 2004), semantic role labelling (Màrquez et al., 2008), cross-linguistic transfer of semantic annotation (Padó and Lapata, 2009; Fung et al., 2007; van der Plas et al., 2011). The knowledge about the likelihood of external causation might be helpful in the task of detecting implicit arguments of verbs and, especially deverbal nouns (Gerber and Chai, 2012; Roth and Frank, 2012). Knowing, for example, that a verb expresses an externally caused event increases the probability of an implicit causer if an explicit causer is not detected in a particular instance of the verb. Our study should

contribute to the development of formal and extensive representations of grammatically relevant semantic properties of verbs, such as Verb Net (Kipper Schuler, 2005) and PropBank (Palmer et al., 2005).

2 External Causation and the Grammar of Language

The distinction between external and internal causation in events described by verbs is introduced by Levin and Rappaport Hovav (1994) to account for the fact that the alternation is blocked in some verbs such as *bloom* in (2). In Levin and Rappaport Hovav's account, verbs which describe externally caused events alternate (1), while verbs which describe internally caused events do not (2).

(2) a. The flowers suddenly bloomed.
      b. * The summer bloomed the flowers.

The main objection to this proposed generalisation is that it does not account for the cross-linguistic variation. Since the distinction concerns the meaning of verbs, one could expect that the verbs which are translations of each other alternate in all languages. This is, however, often not true. There are many verbs that do alternate in some languages, while their counterparts in other languages do not (Alexiadou et al., 2006; Schafer, 2009; Alexiadou, 2010). For example, *appear* and *arrive* do not alternate in English, but their equivalents in Japanese or in the Salish languages do.

To account for the variation in cross-linguistic data Alexiadou (2010) introduces the notion of cause-unspecified events, a category between externally caused and internally caused events. Introducing gradience into the classification allows Alexiadou to propose generalisations which apply across languages: cause-unspecified verbs alternate in all languages, while only some languages allow the alternation if the event is either externally or internally caused. To allow the alternation in the latter cases, languages need a special grammatical mechanism. In English, for example, this mechanism is not available, which is why only cause-unspecified verbs alternate. The alternation is thus blocked in both verbs describing externally caused and internally caused events.

Alexiadou's account is based not only on the observations about the availability of the alternation, but also about morphological encoding of the alternation across languages. Unlike English,

which does not mark the alternation morphologically (note that the two versions of English verbs in (1-3) are morphologically identical), other languages encode the alternation in different ways, as shown in (3).

		Causative	Anticausative
	Mongolian	xajl-**uul**-ax	xajl-ax
		'melt'	'melt'
(3)	Russian	rasplavit	rasplavit-**sja**
		'melt'	'melt'
	Japanese	atum-**eru**	atum-**aru**
		'gather'	'gather'

An analysis of the distribution of morphological marking across languages leads Haspelmath (1993) to introduce the notion of likelihood into his account of the meaning of the alternating verbs. In a study of 31 verbs in 21 languages from all over the world, Haspelmath notices that certain verbs tend to get the same kind of marking across languages. For each verb, he calculates the ratio between the number of languages which mark the anticausative version and the number of languages which mark the causative version of the verb. He interprets this ratio as a quantitative measure of how spontaneous events described by the verbs are. As each verb is assigned a different score, ranking the verbs according to the score results in a "scale of increasing likelihood of spontaneous occurrence". Events with a low anticausative/causative ratio (e.g. *boil*, *dry*, *melt*) are likely to occur spontaneously, while events with a high ratio (e.g. *break*, *close*, *split*) are likely to be caused by an external causer.

3 The Model

Our study pursues the quantitative assessment of the likelihood of external causation in the events described the alternating verbs. We estimate the likelihood by means of a Bayesian model which divides events into classes based on the distribution of causative and anticausative uses of verbs in a corpus. By varying the settings of the model, we address two questions discussed in the linguistic literature: 1) Is the distinction between externally caused and internally caused events binary,, as argued by Levin and Rappaport Hovav (1994), or are there are intermediate classes, as argued by Alexiadou (2010)? and 2) Do we obtain better estimation of the likelihood from cross-linguistic than from monolingual data?

We design a probabilistic model which estimates the likelihood of external causation and generates a probability distribution over a given number of event classes for each verb in a given set of verbs. The model formalises the intuition that an externally caused event tends to be expressed by a verb in its causative realisation. In other words, if the likelihood of external causation of the event is encoded in the use of the verb which describes the event, then the causer is expected to appear frequently in the realisations of the verb. The opposite is expected for internally caused events. Cause-unspecified events are expected to appear with and without the causer equally.

To take into account the two questions discussed in the theoretical approaches, namely the number of classes and the role of cross-linguistic data in the classification of events, we design four versions of the model, varying the input data and the number of classes in the output: a) monolingual input and two classes; b) cross-linguistic input and two classes; c) monolingual input and three classes; d) cross-linguistic input and three classes.

The current cross-linguistic versions of the model include only two languages, English and German, because we test the models in a minimal cross-linguistic setting. In principle, the approach can be easily extended to include any number of languages.

As it can be seen in its graphical representation in Figure 1, the model consists of three variables in the monolingual version and of four variables in the cross-linguistic version.
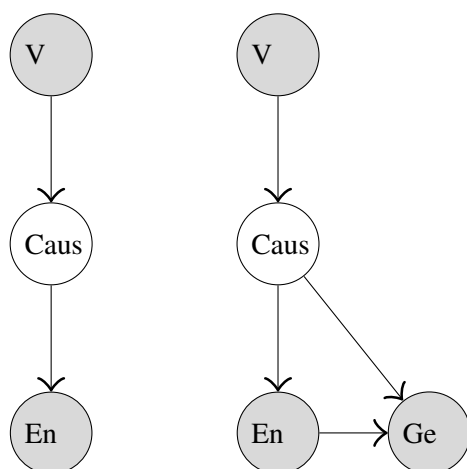
The first variable in both versions is the set of verbs V. This can be any given set of verbs.

The second variable is the event class of the verb, for which we use the symbol $Caus$. The values of this variable depend on the assumed classification. In the two-class version, the values are *causative*, representing externally caused events, and *anticausative*, representing internally caused events. In the three-class version, the variable can take one more value, *unspecified*, representing cause-unspecified events.

The third (En) and the fourth (Ge) (in the cross-linguistic version) variables are the surface realisations of the verbs in parallel instances. These variables take three values: *causative* for active transitive use, *anticausative* for intransitive use, and *passive* for passive use in the languages in question.

We represent the relations between the variables as a Bayesian network. The variable that represents the event class of verbs ($Caus$) is unobserved. The values for the other three variables are observed in the data source. Note that the input to the model does not contain the information about the event class at any point.

The dependence between En and Ge in the bilingual version of the model represents the fact that the two instances of a verb are translations of each other, but does not represent the direction of translation in the actual data. The form of the instance in one language depends on the form of the parallel instance because they express the same meaning in the same context, regardless of the direction of translation.

Assuming that the variables are related as in Figure 1, En and Ge are conditionally independent of V given $Caus$, so we can calculate the probability of the model as in (4) for the monolingual version and as in (6) for the cross-linguistic version.

$$P(v, caus, en) = P(v) \cdot P(caus|v) \cdot P(en|caus) \tag{4}$$

$$P(caus|v, en) = \frac{P(v) \cdot P(caus|v) \cdot P(en|caus)}{\sum_{caus} P(v) \cdot P(caus|v) \cdot P(en|caus)} \tag{5}$$

We estimate the conditional probability of the event class given the verb ($P(caus|v)$) by querying the model, as shown in (5) for the monolingual version and in (7) for the bilingual version..

Figure 1: Two version of the Bayesian net model for estimating external causation.

$$P(v, caus, en, ge) = \tag{6}$$
$$P(v) \cdot P(caus|v) \cdot P(en|caus) \cdot P(ge|caus, en)$$

$$P(caus|v, en, ge) = \tag{7}$$
$$\frac{P(v) \cdot P(caus|v) \cdot P(en|caus) \cdot P(ge|caus, en)}{\sum_{caus} P(v) \cdot P(caus|v) \cdot P(en|caus) \cdot P(ge|caus, en)}$$

We assign to each verb the event class that is most probable given the verb, as in (8).

$$caus\_class(verb) = arg\,max_{caus} P(caus|v) \tag{8}$$

All the variables in the model are defined so that the parameters can be estimated on the basis of frequencies of instances of verbs automatically extracted from parsed corpora.

4 Experiments

The accuracy of the predictions of the model is evaluated in experiments.

4.1 Materials and Methods

The verbs for which we estimate the likelihood are the 354 verbs that participate in the causative alternation in English, as listed by Levin (1993), and the 26 verbs listed as alternating in a typological study (Haspelmath, 1993).

We estimate the parameters of the model by implementing the expectation-maximisation algorithm. The algorithm is initialised by assigning different arbitrary values to the parameters of the model. The classification reported in the paper is obtained after 100 iterations.

We train the classifier using the data automatically extracted from an English-German parallel corpus (Europarl (Koehn, 2005)). Both monolingual and bilingual input data are extracted from the parallel corpus. All German verbs which are word-aligned with the alternating English verbs listed in the literature are regarded as German equivalents. By extracting cross-linguistic equivalents automatically from a parallel corpus, we avoid manual translation into German of the lists of English verbs discussed in the literature. In this way, we eliminate the judgements which would be involved in the process of translation.

The corpus is syntactically parsed (using the MaltParser (Nivre et al., 2007)) and word-aligned

(using GIZA++ (Och and Ney, 2003)). For both the syntactic parses and word alignments, we reuse the data provided by Bouma et al. (2010).

We extract only the instances of verbs where both the object (if there is one) and the subject are realised in the same clause, excluding the instances involving syntactic movements and coreference. Transitive instances are considered causative realisations, intransitive anticausative. We count passive instances separately because they are formally transitive, but they usually do not express the causer.

German equivalents of English alternating verbs are extracted in two steps. First, all verbs occurring as transitive, intransitive, and passive were extracted from the German sentences that are sentence-aligned with the English sentences containing the instances of alternating verbs. These instances were considered candidate translations. The instances that are the translations of the English instances were then selected on the basis of word alignments. Instances where at least one element (the verb, the head of its object, or the head of its subject) is aligned with at least one element in the English instance were considered aligned.

Only the instances of English verbs that are translated with a corresponding finite verbal form in German are extracted, excluding the cases where English verbs are translated into a corresponding non-finite form such as infinitive, nominalization, or participle in German.

4.2 Evaluation

We evaluate the performance of the models against the scale of spontaneous occurrence proposed by Haspelmath (1993), shown in (9). We expect the verbs classified as internally caused by our models to correspond to the verbs with a low morphological anticausative/causative ratio (those on the left side of the scale). The opposite is expected for externally caused verbs. Cause-unspecified verbs are expected to be in the middle of Haspelmath's scale.

(9) *boil, dry, wake up, sink, learn-teach, melt,*
stop, turn, dissolve, burn, fill, finish, begin,
spread, roll, develop, rise-raise, improve,
rock, connect, change, gather, open, break,
close, split

To evaluate the output of our models against the scale, we discretise the scale so that the agreement

is maximised for each version of the model. For example, the threshold which divides the verbs into anticausative and causative in the two-way classification is set after the verb *turn*.

By evaluating the performance of our models against a typology-based measure, we avoid eliciting human judgements, which is a known problem in computational approaches to causality. The downside of this approach is that such evaluation is currently possible for a relatively small number of verbs.

5 Results and Discussion

Table 1 shows all the confusion matrices of the classifications performed automatically in comparison with the classifications based on the typology rankings.[1]

In the two-way classification, the two versions of the model, with monolingual and with bilingual input, result in identical classifications. The agreement of the models with the typological ranking can be considered very good (85%). The optimal threshold divides the verbs into two asymmetric classes: eight verbs in the internally caused class and eighteen in the externally caused class. The agreement is better for the internally caused class.

In the three way-classification, the performance of both versions of the model drops. In this setting, the output of the two versions differs: there are two verbs which are classified as externally caused by the monolingual version and as cause-unspecified by the bilingual version, which results in a slightly better performance of the bilingual version. Given the small number of evaluated verbs, however, this tendency cannot be considered significant.

The three-way classification is more difficult than the two-way classification, but the difficulty is not only due to the number of classes, but also to the fact that two of the classes are not well-distinguished in the data. While the class of internally caused events is relatively easily distinguished (small number of errors in all classifications), the classes of externally caused and cause-unspecified verbs are hard to distinguish. This finding supports the two-way classification argued for in the literature.

The classification performed by the bilingual

model indicates that the distinction between externally caused and cause-unspecified verbs might still exist. Compared to the monolingual classification, more verbs are classified as cause-unspecified, and they are grouped in the middle of the typological scale. Since the model takes into account cross-linguistic variation in the realisations of the alternating verbs, the observed difference in the performance could be interpreted as a sign that the distinction between cause-unspecified and externally caused events does emerge in cross-linguistic contexts.

While supporting the two-way classification of events, our experiments do not provide a definite answer to the question of whether there are more than two classes of events. To obtain significant results, more verbs need to be evaluated. However, the typological data used in our experiments (Haspelmath, 1993) are not easily available. This kind of data are currently not included in the typological resources (such as the WALS database (Dryer and Haspelmath, 2013)), but they can, in principle, be collected from other electronic sources of language documentation, which are increasingly available for many different languages.

6 Related Work

The proposed distinction between externally and internally caused events is addressed by McKoon and Macfarland (2000). They study twenty-one verbs defined in the linguistic literature as describing internally caused events and fourteen verbs describing externally caused events. Their corpus study shows that the appearance of these verbs as causative (transitive) and anticausative (intransitive) cannot be used as a diagnostic for the kind of meaning that has been attributed to them.

Since internally caused verbs do not enter the alternation, they were expected to be found in intransitive clauses only. This, however, was not the case. The probability for some of these verbs to occur in a transitive clause is actually quite high (0.63 for the verb *corrode*, for example). More importantly, no difference was found in the probability of the verbs denoting internally caused and externally caused events to occur as transitive or as intransitive. This means that the acceptability judgements used in the qualitative analysis do not apply to all the verbs in question, and, also, not to all the instances of these verbs.

---

[1]The table contains 26 instead of 31 verbs because corpus data could not be reliably extracted for some phrasal verbs listed by Haspelmath.

Model	2-class				3-class					
	Monolingual		Bilingual		Monolingual			Bilingual		
Typology	acaus	caus	acaus	caus	acaus	caus	unspec.	acaus	caus	unspec.
acaus	**8**	0	**8**	0	**6**	0	1	**6**	0	1
caus	4	**14**	4	**14**	0	**3**	0	0	**3**	0
unspec.	—	—	—	—	4	5	**7**	4	3	**9**
Agreement	85%		85%		61%			69%		

Table 1: Per class and overall agreement between the corpus-based and the typology-based classification of verbs; acaus = internally caused, caus = externally caused, unspec. = cause-unspecified.

Even though the most obvious prediction concerning the corpus instances of the two groups of verbs was not confirmed, the corpus data were still found to support the distinction between the two groups. Examining 50 randomly selected instances of transitive uses of each of the studied verbs, McKoon and Macfarland (2000) find that, when used in a transitive clause, internally caused change-of-state verbs tend to occur with a limited set of subjects, while externally caused verbs can occur with a wider range of subjects. This difference is statistically significant.

The relation between frequencies of certain uses and the lexical semantics of English verbs has been explored by Merlo and Stevenson (2001) in the context of automatic verb classification. Merlo and Stevenson (2001) show that information collected from instances of verbs in a corpus can be used to distinguish between three different classes which all include verbs that alternate between transitive and intransitive use. The classes in question are manner of motion verbs (10), which alternate only in a limited number of languages, externally caused change of state verbs (11), alternating across languages, and performance/creation verbs, which are not lexical causatives (12).

(10) a. The horse raced past the barn.
    b. The jockey raced the horse past the barn.

(11) a. The butter melted in the pan.
    b. The cook melted the butter in the pan.

(12) a. The boy played.
    b. The boy played soccer.

One of the most useful features for the classification proved to be the *causativity* feature. It represents the fact that, in the causative alternation, the same lexical items can occur both as subjects and as objects of the same verb. This feature sets apart the two causative classes from the performance class.

In the context of psycholinguistic empirical approaches to encoding causality in verbs, it has been established that assigning a causal relation to a sequence of events can be influenced by the native languages (Wolff et al., 2009a; Wolff and Ventura, 2009b). English speakers, for instance, tend to assign causal relations more than Russian speakers.

In our study, we draw on the fact that the semantic properties of verbs are reflected in the way they are used in a corpus, established by the previous studies. We explore this relationship further, relating it to a deeper semantic analysis and to the typological distribution of grammatical features.

7 Conclusion and Future Work

The experiments presented in this article provide empirical evidence that contribute to a better understanding of the relationship between the causal semantics of verbs, their formal morphological and syntactic properties, and the variation in their use. We have shown that the likelihood of external causation of events is encoded in the distribution of the causative and anticausative uses of verbs. Two classes of events, externally caused and internally caused events, can be distinguished automatically based on corpus data.

In future work, we will further investigate the question of whether there are more than two classes of events and how they are distinguished. We will explore potential tendencies indicated by our findings. We will apply the approach proposed in this article to an extended data set. On one hand, we will collect typological data for more verbs, exploring possibilities of automatic data extraction. On the other hand, we will include more languages in the model to ensure a better representation of cross-linguistic variation.

References

Artemis Alexiadou, Elena Anagnostopoulou, and Florian Schfer. 2006. The properties of anticausatives crosslinguistically. In Mara Frascarelli, editor, *Phases of Interpretation*, pages 187–212, Berlin, New York. Mouton de Gruyter.

Artemis Alexiadou. 2010. On the morpho-syntax of (anti-)causative verbs. In Malka Rappaport Hovav, Edit Doron, and Ivy Sichel, editors, *Syntax, Lexical Semantics and Event Structure*, pages 177–203, Oxford. Oxford University Press.

Steven Bethard. 2007. *Finding Event, Temporal and Causal Structure in Text: A Machine Learning Approach*. Ph.D. thesis, University of Colorado at Boulder.

Gerlof Bouma, Lilja Øvrelid, and Jonas Kuhn. 2010. Towards a large parallel corpus of cleft constructions. In *Proceedings of the Seventh conference on International Language Resources and Evaluation (LREC'10)*, Valletta, Malta. European Language Resources Association.

Matthew S. Dryer and Martin Haspelmath, editors. 2013. *WALS Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig.

Pascale Fung, Zhaojun Wu, Yongsheng Yang, and Dekai Wu. 2007. Learning bilingual semantic frames: Shallow semantic parsing vs. semantic role projection. In *11th Conference on Theoretical and Methodological Issues in Machine Translation (TMI 2007)*, pages 75–84, Skovde, Sweden.

Matthew Gerber and Joyce Y. Chai. 2012. Semantic role labeling of implicit arguments for nominal predicates. *Computational Linguistics*, 38(4):755–798.

Thomas Gilovich, Robert Vallone, and Amos Tversky. 1985. The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology*, 17(3):295–314.

Cécile Grivaz. 2012. *Automatic extraction of causal knowledge from natural language texts*. Ph.D. thesis, University of Geneva.

Martin Haspelmath. 1993. More on typology of inchoative/causative verb alternations. In Bernard Comrie and Maria Polinsky, editors, *Causatives and transitivity*, volume 23, pages 87–121, Amsterdam/Philadelphia. John Benjamins Publishing Co.

Karin Kipper Schuler. 2005. *VerbNet: A broad-coverage, comprehensive verb lexicon*. Ph.D. thesis, University of Pennsylvania.

Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *Proceedings of MT Summit 2005*, Phuket, Thailand.

Mirella Lapata and Chris Brew. 2004. Verb class disambiguation using informative priors. *Computational Linguistics*, 30(1):45–73.

Beth Levin and Malka Rappaport Hovav. 1994. A preliminary analysis of causative verbs in English. *Lingua*, 92:35–77.

Beth Levin and Malka Rappaport Hovav. 2005. *Argument realization*. Cambridge University Press, Cambridge.

Beth Levin. 1993. *English verb classes and alternations : a preliminary investigation*. The University of Chicago Press, Chicago.

Lluís Màrquez, Xavier Carreras, Kenneth C. Litkowski, and Suzanne Stevenson. 2008. Semantic role labeling: An introduction to the special issue. *Computational Linguistics*, 34(2):145–159.

Gail McKoon and Talke Macfarland. 2000. Externally and internally caused change of state verbs. *Language*, 76(4):833–858.

Paola Merlo and Susanne Stevenson. 2001. Automatic verb classification based on statistical distribution of argument structure. *Computational Linguistics*, 27(3):373–408.

Albert Michotte. 1963. *The perception of causality*. Basic Books, Oxford, England.

Jacques Moeschler. 2011. Causal, inferential and temporal connectives: Why 'parce que' is the only causal connective in French. In S. Hancil, editor, *Marqueurs discursifs et subjectivité*, pages 97–114, Rouen. Presses Universitaires de Rouen et du Havre.

Joakim Nivre, Johan Hall, Jens Nilsson, Chanev Atanas, Gleşen Eryiğit, Sandra Kbler, Svetoslav Marinov, and Erwin Marsi. 2007. MaltParser: A language-independent system for data-driven dependency parsing. *Natural Language Engineering*, 13(2):95–135.

Franz Josef Och and Hermann Ney. 2003. A systematic comparison of various statistical alignment models. *Computational Linguistics*, 29(1):19–52.

Sebastian Padó and Mirella Lapata. 2009. Cross-lingual annotation projection of semantic roles. *Journal of Artificial Intelligence Research*, 36:307–340.

Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The Proposition Bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31(1):71–105.

James Pustejovsky. 1995. *The generative lexicon*. MIT Press, Cambridge, MA.

Gillian Ramchand. 2008. *Verb Meaning and the Lexicon: A First Phase Syntax*. Cambridge Studies in Linguistics. Cambridge University Press, Cambridge.

Michael Roth and Anette Frank. 2012. Aligning predicate argument structures in monolingual comparable texts: A new corpus for a new task. In *SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pages 218–227, Montréal, Canada, 7-8 June. Association for Computational Linguistics.

Florian Schafer. 2009. The causative alternation. In *Language and Linguistics Compass*, volume 3, pages 641–681. Blackwell Publishing.

Leonard Talmy. 2000. *Towards a cognitive semantics*. The MIT Press, Cambridge Mass.

Amos Tversky and Daniel Kahneman. 1982. Causal schemata in judgments under uncertainty. In Daniel Kahneman, Paul Slovic, and Amos Tversky, editors, *Judgement Under Uncertainty: Heuristics and Biases*.

Lonneke van der Plas, Paola Merlo, and James Henderson. 2011. Scaling up automatic cross-lingual semantic role annotation. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 299–304, Portland, Oregon, USA, June. Association for Computational Linguistics.

Phillip Wolff and Tatyana Ventura. 2009b. When Russians learn English: How the semantics of causation may change. *Bilingualism: Language and Cognition*, 12(2):153–176.

Phillip Wolff, Ga-Hyun Jeon, and Yu Li. 2009a. Causal agents in English, Korean and Chinese: The role of internal and external causation. *Language and Cognition*, 1(2):165–194.

Sandrine Zufferey. 2012. 'Car, parce que, puisque' revisited: Three empirical studies on French causal connectives. *Journal of Pragmatics*, 44:138–153.