

A general framework for the annotation of causality based on FrameNet

Laure Vieu[†], Philippe Muller[†], Marie Candito^{*}, Marianne Djemaa^{*}

[†]IRIT CNRS & Toulouse University, ^{*}Alpage INRIA & Paris Diderot

Abstract

We present here a general set of semantic frames to annotate causal expressions, with a rich lexicon in French and an annotated corpus of about 4000 instances of causal lexical items with their corresponding semantic frames. The aim of our project is to have both the largest possible coverage of causal phenomena in French, across all parts of speech, and have it linked to a general semantic framework such as FN, to benefit in particular from the relations between other semantic frames, e.g., temporal ones or intentional ones, and the underlying upper lexical ontology that enables some forms of reasoning. This is part of the larger ASFALDA French FrameNet project, which focuses on a few different notional domains which are interesting in their own right (Djemaa et al., 2016), including cognitive positions and communication frames. In the process of building the French lexicon and preparing the annotation of the corpus, we had to remodel some of the frames proposed in FN based on English data, with hopefully more precise frame definitions to facilitate human annotation. This includes semantic clarifications of frames and frame elements, redundancy elimination, and added coverage. The result is arguably a significant improvement of the treatment of causality in FN itself.

Keywords: semantic role labelling, causality, FrameNet

1. Introduction

A lot of information in natural language is of a causal nature: relations between events, explanations, argumentations are all important to the understanding of texts, and thus useful in question-answering (Oh et al., 2013), information extraction (Khoo et al., 2000; nan Cao et al., 2014) or textual entailment (Gordon et al., 2012). Davidson, in his work on individuation of events (Davidson, 1969), went as far as saying causality defines events, as “Events are identical if and only if they have exactly the same causes and effects”, and causality is an active area of research in linguistics (Neeleman and van de Koot, 2012). Causality is also an important notion in discourse analysis, being one of the most important categories of discourse relations (Hovy and Maier, 1992; Mak and Sanders, 2013).

Numerous studies in NLP have focused on extraction of cause-effect relationships between nouns denoting events (Kozareva, 2012), or relations between events denoted by verbs (Riaz and Girju, 2013; Do et al., 2011), or combinations of nouns and verbs (Riaz and Girju, 2014). Other researchers try instead to gather general knowledge about typical causal links, mostly expressed by verbs (Hashimoto et al., 2009; Szpektor and Dagan, 2008; Chklovski and Pantel, 2004), essentially for textual entailment or knowledge mining.

They all target realized causalities: an event occurred and caused another event, or is likely to be the cause of another event. This is also the focus of annotation efforts, integrating time and causal information, as in (Mirza and Tonelli, 2014; Bethard and Martin, 2008). They make distinctions between groups of verbs expressing such causalities, based on the work of (Wolff and Song, 2003), namely *cause*, *enable*, *prevent*. Other projects aim at annotating discourse relations between clauses, among which causal relations, marked or not by discourse connectives such as *because* or *then* (Prasad et al., 2008; Carlson et al., 2007).

In contrast, larger lexical projects have on the one hand covered all sorts of POS expressing causality (including nouns, verbs, adverbs, conjunctions, prepositions, adjectives) and on the other hand, distinguished a much larger set of causal

relationships involving events as well as facts: for instance in FrameNet (henceforth FN) (Baker et al., 1998), some frames are concerned with argumentation, where typical causal expressions introduce evidence for a claim, or reasons for an agent’s behaviour.

This is consistent with more psycholinguistically oriented studies, which clearly distinguish the expression of factual and epistemic causality (Spooren et al., 2010).

Examples in (1), taken from FN, show relationships between “situations”, events or facts, with names specific to each frame, and which parallel cause-effect relationships (in bold, the lexical unit triggering the frame annotation, called a frame-evoking element in FN).¹

- (1) a. [What happened to these two chaps]_{support}
proves [the rumour is not true]_{proposition}·
EVIDENCE
- b. I like to think the main **reason** [we]_{Agent} [have stayed together since the World Cup]_{action} is [the great spirit within this Australian team]_{state_of_affairs}· REASON
- c. If [such a small earthquake]_{cause} **causes** [problems]_{effect}, just imagine a big one!
CAUSATION

Similar distinctions appear in (Dunietz et al., 2015), who annotated four kinds of causalities: *consequence*, which seems similar to the FN CAUSATION frame, *motivation* (similar to REASON), *purpose* (similar to FN’s PURPOSE), and *inference* (similar to EVIDENCE). Unfortunately their effort was made without an explicit relation to an existing semantic lexicon, and has only a small dataset to support the framework. To our knowledge, FN is the semantic lexicon which has the largest coverage and the most detailed analysis of the expression of causality in English, equipped with annotation procedures.² The detailed analysis relies

¹In the remaining of the paper, Frame names are written in small caps for clarity.

²Annotation in FN is nevertheless deliberately limited to the sentence, and as a result only partially accounts for causal relations between clauses across different sentences: only the argu-

on the distinction of a large number of frames as well as the identification of various “frame elements” in each frame. Frame element is the FN term for a frame-specific semantic role, but in the following we will use the shorter term “role”, and will call “role filler” the linguistic material that expresses a given role.

We present here a general set of semantic frames to annotate causal expressions, with a rich lexicon in French and an annotated corpus of about 4000 instances of causal lexical items with their corresponding semantic frames. The aim of our project is to have both the largest possible coverage of causal phenomena in French, across all parts of speech, and have it linked to a general semantic framework such as FN, to benefit in particular from the relations between other semantic frames, e.g., temporal ones or intentional ones, and the underlying upper lexical ontology that enables some forms of reasoning. This is part of the larger ASFALDA French FrameNet project, which focuses on a few different domains which are interesting in their own right (Candito et al., 2014; Djemaa et al., 2016), including cognitive stances and communication frames.

In the process of building the French lexicon and preparing the annotation of the corpus, we had to remodel some of the frames proposed in FN based on English data, with hopefully more precise frame definitions to facilitate human annotation. This includes semantic clarifications of frames and roles, redundancy elimination, and added coverage. The result is arguably a significant improvement of the treatment of causality in FN itself.

Section 2. describes the causality-related frames used. Section 3. presents the methodology we used, from building the lexicon to the corpus annotation effort, including an evaluation of the inter-annotator agreement, and some statistics on the resulting annotated resource.

2. Causality related frames

To identify the relevant causal frames in FN, we followed the procedure described in (Candito et al., 2014). We used the 1.5 FrameNet release. Exploring the set of frames and frame-to-frame relations (e.g., inheritance) we gathered 14 general frames expressing causality : CAUSATION, CAUSE_TO_START, CONTINGENCY, EVIDENCE, EXPLAINING_THE_FACTS, INTENTIONALLY_AFFECT, LAUNCH_PROCESS, OBJECTIVE_INFLUENCE, MAKE_POSSIBLE_TO_DO, PREVENTING, REASON, RESPONSE, SIGN, TRANSITIVE_ACTION , of which TRANSITIVE_ACTION is a high-level non lexicalized frame, and INTENTIONALLY_AFFECT lexicalized by very few and general lexical units only (e.g. *to do* in *to do the dishes*), so in the remainder we focus on the 12 lexicalized frames. Many more specialized frames do involve causality, most of them subsumed by TRANSITIVE_ACTION, like for instance KILLING or RELEASING and the 32 ones named CAUSE_XX or CAUSE_TO_YY (e.g., CAUSE_EXPANSION, CAUSE_TO_PERCEIVE). We chose not to consider them in

ments of the causal relation appearing in the same sentence as the causal lexical unit are annotated. The other ones are not recoverable. Frames are only evoked by lexical units, so annotation of unmarked causal relations between clauses is also left out. But both these limitations may be lifted in an extension.

this work, for being too specific. A CAUSATION_SCENARIO was also identified, not lexicalized, and not considered further.

A contrastive analysis of the semantics of these remaining 12 general lexicalized frames, exploiting existing distinctions in the literature on causality, (Zufferey and Cartoni, 2012; Atallah, 2014; Degand and Pander Maat, 2003) and a careful examination of the boundaries with frames belonging to other domains such as the cognitive stances domain prompted a clarification of the frames’ characterization and of their roles. This process brought us to merge some frames (noted XX+YY), to eliminate one frame (SIGN, for being redundant with EVIDENCE), to add new ones (CAUSE_ENUNCIATION and ATTRIBUTING_CAUSE), and to adjust the set of core roles in a few cases. The objective in this was to be able to better carve the boundaries between frames for polysemous lexical units and obtain a better coverage of the French lexicon. This resulted in a remodeled set of 11 lexicalized frames, which we will now briefly describe. These 11 general lexicalized frames can be grouped by general semantic and pragmatic features.

Causation and its subtypes Standard “objective” causation between two events or situations, a cause and an effect, is accounted for by the frame CAUSATION. Cause and effect can be presented in the iconic way or in a backward way, viz. *The power grid failed so the lights went off* / *The lights went off because the power grid failed*, and the cause event may be explicit or the implicit action of an ‘actor’ viz. *the attacks caused considerable damage* / *the soldiers caused considerable damage*. Two variants of standard causation, CAUSE_TO_START+LAUNCH_PROCESS and EXPLAINING_THE_FACTS, are distinguished for their particular presentation of the facts: the effect having just started for the former (e.g. *trigger a crisis*), or the effect being already known to have occurred and its origin debated for the latter (e.g. *the present crisis explains our problems*).

The frame REASON also is a special case of CAUSATION (although not originally related with frame inheritance) in which the effect is the action of an agent, and the causation actually expressed is between the cause (or, again, an actor of an implicit cause) and the decision of the agent to do the effect, something sometimes called volitional causation (Degand and Pander Maat, 2003). RESPONSE is itself a special case of REASON in which the cause is the action of another sentient and the effect is described as a reaction to that action³.

In all these cases, distinctions between a frame and its subtypes are done directly on the lexicon, not in context, and no polysemy is ever introduced on a frame and its subtypes (they do not share frame-evoking elements).

- (2) [Ce transfert]_{Action} était **motivé** par [la crainte d’une OPA]_{State of affairs} (REASON)
(*This transfer was motivated by the fear of a takeover.*)
- (3) [Les constructeurs]_{Agent} ont **réagi** [cette phase de

³Examples in French in the following are either taken from our corpus or slightly adapted for brevity.

ralentissement des ventes]Trigger en [reprenant l'offensive]Response (RESPONSE)
(Manufacturers reacted to this slowing-down of sales by taking up the initiative.)

Finally, PREVENTING, while not a proper subtype of causation, is the causation of an effect in the form of a negated event or situation.

Evidence Some degree of argumentation is already implicit in REASON and RESPONSE, as the agent exploited the cause to motivate its action. In the terms of Degand & Pander Maat (Degand and Pander Maat, 2003), we are on a higher scale of “speaker involvement” (this term actually covers intentions and other attitudes of any sentient, not just the speaker). Argumentation proper occurs when a speaker presents a (less debatable) event or fact as support for a debatable claim or proposition. It is accounted for by the frame EVIDENCE. As noted in studies on “subjective” or epistemic causal relations (Sweetser, 1990; Degand and Pander Maat, 2003; Sanders and Stukker, 2012) what is presented as support is the cause of getting to know or being more convinced of the proposition. Since the proposition may be the factual cause of the support, as in the classic abduction example *He is away, because the lights are out* involving a discourse marker as well as in *The proliferation of disputes over scarce water resources (...) at-tests to the impact of resource depletion*, EVIDENCE is not a special case of CAUSATION. EVIDENCE has been adjusted in various ways. For instance, several cases covered in EVIDENCE are instead considered as cases of PROVING (a new frame created in the distinct cognitive stance domain), on the grounds that the claim’s support is absent as in *Earlier chapters have demonstrated this assessment* (the support is described in the chapters). On the other hand, EVIDENCE now takes over some cases left by the removal of the frame SIGN (the rest being redirected to the frame OMEN, not in the causal domain).

Cause and speech acts The new frame CAUSE_ENUNCIATION has been introduced for cases of speech act causation, well distinguished from epistemic causation (Sweetser, 1990), but absent from the original FN. In the classic example *Are you ready? Because we are late*, the being-late situation causes the interrogative speech act, not the state of being ready or not. This is quite rare in our corpus as it is much more frequent in speech than in written texts, but still occurs, especially under a specific case of explaining an implicature or a presupposition, as in *John came with his sister, because he has a sister*.

A new frame ATTRIBUTING_CAUSE has also been introduced involving causality and communication, to distinguish between an explanation and a reported explanation, as in *The crisis explains the state of the economy* (EXPLAINING_THE_FACTS) versus *Officials explain the crisis by poor management*. (ATTRIBUTING_CAUSE).

Objective influence CONTINGENCY+OBJECTIVE_INFLUENCE characterizes the influence or dependence of a process on another or between two parameters of a process, without the temporal precedence typical of causation between events.

(4) La [situation]dependent restait très confuse, alimentée par de [multiples rumeurs et déclarations];influencer.
(The situation remained confused, fueled (lit. fed) by rumours and declarations)

MAKE_POSSIBLE_TO_DO is a subcase of OBJECTIVE_INFLUENCE in the original FrameNet hierarchy, in which a situation or agent (the enablement) helps bring about an event (the enabled action).

(5) [Nos industries ont amélioré leur compétitivité];enablement, ce [qui];enablement a autorisé [une forte croissance de nos parts de marché];enabled_action.
(Our industries improved their competitiveness, which allowed for a strong growth in market shares.)

Let’s observe that the frame PURPOSE, a relation often considered as causal is here deemed to belong to the cognitive stances domain as we have limited the causality domain to effective causal relations and the purpose motivating an action may never be reached.

Figure 1 presents all considered frame and their inheritance relations in FrameNet.

3. Methodology

Most of previous works on annotating causality have adopted a text-driven methodology, where a corpus is selected, and every causal expression or relation corresponding to the scope of the given project is annotated (Bethard and Martin, 2008; Do et al., 2011; Mirza and Tonelli, 2014). In contrast, the work of (Dunietz et al., 2015) started from a lexicon of expressions related to causality before giving full-text annotation to annotators, showing this improves inter-annotator agreement. This is in fact the methodology followed in a broader semantic project such as FN, where the annotation is done frame by frame, each frame being defined beforehand along with its lexical units, i.e. the lexical items that may evoke a frame, hereafter LUs. FN then selects various disambiguated occurrences of LUs. This is supposed to ensure full lexical coverage for a given frame, and a maximal variety of syntactic constructions, although it biases the annotations towards the senses of the covered frames (which is not a problem when a focus on a few frames is chosen as for causality) and does not preserve the naturally occurring frequencies of syntactic realizations of roles. FN later added full-text annotations to have a more representative coverage.

We followed a similar path for the project, but aiming for a complete lexical coverage for causality: we first defined a lexicon of causal lexical units for French beforehand. The subsequent corpus annotation phase, performed within the ASFALDA French FrameNet project is described at length in (Djemaa et al., 2016) and concerns three other notional domains, on top of causality. We used two syntactically annotated corpus, the French Treebank (Abeillé and Barrier, 2004) and the Sequoia treebank (Candito and Seddah, 2012), in which we isolated the occurrences of the lemmas of our causal LUs and had them disambiguated and

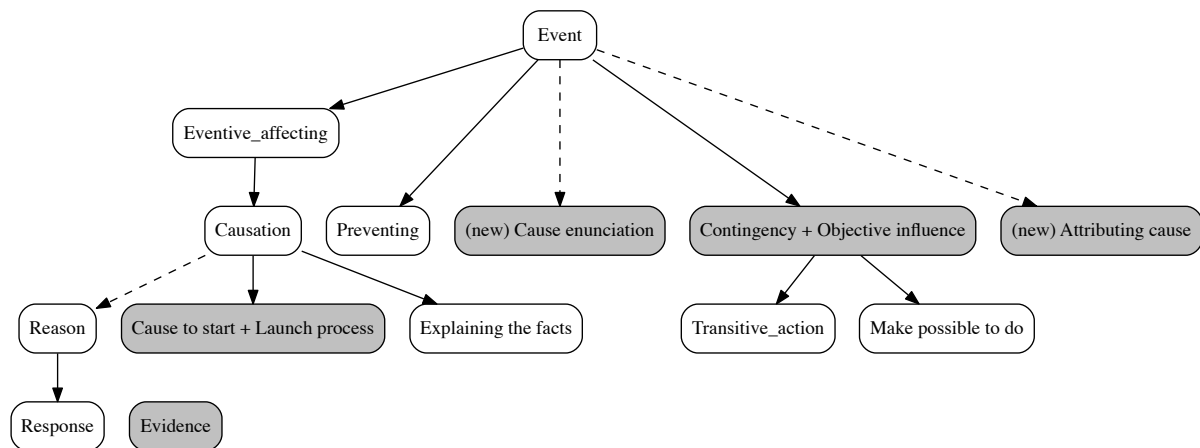


Figure 1: Causal frame hierarchy; Grey nodes indicate frame changes from Framenet 1.5, either new, merged or redefined frames. Dotted relations are additions. Transitive action is an unlexicalised parent to more specific predicates. NB: Evidence is not related to FrameNet hierarchy.

annotated according to our French FN frames. Annotators had to disambiguate between causal senses and frames selected from other domains within the ASFALDA French FrameNet project, or “other” for senses outside the perimeter of the project. This methodology preserves the natural sense and role-realization probabilistic distributions, modulo an upper bound of 100 occurrences on frequent LUs (see section 3.2.).

3.1. Lexicon

The lexicon was built using multiple lexicographic sources: translation of the English lexical units given automatically from a parallel corpus from (Mouton et al., 2010; Padó, 2007), extended with French synonymy resources, the causal items from a lexicon of French discourse connectives (Roze et al., 2012) and input from linguists within the project, and then manually validated. The manual validation was done by two members of the project separately then adjudicated. Overall, 332 different lexical units were listed for causal frames.

3.2. Frame annotation

For frame annotation we used the Salto tool (Burchardt et al., 2006), developed by the German FrameNet project Salsa. Salto allows for the annotation of frames evoked by lexical units, and which linguistic material fills which role (the role fillers) on top of a syntactic annotation (in our case provided by the syntactic treebanks). We preannotated the corpus for each occurrence of every item listed as a potential causality LU with the frames associated to the LU in the lexicon, and then extracted the relevant sentences (more precisely, for each causality LU listed in our lexicon, we limited ourselves to the first 100 occurrences at most). For each selected occurrence, two annotators had to independently decide whether the occurrence corresponded to one of the proposed frames and if so annotate their core role fillers, otherwise a special Other_sense frame was to be

used.

The annotation guide was written by experts for each domain (the present authors), after a sample of annotations, and each potentially ambiguous lexical item was analyzed with typical constructions. It was updated during the campaign by taking into account annotators’ feedback. A couple of disambiguation principles are worth commenting.

First, complying with the lexical orientation of FN, frames related by inheritance are assumed not to give rise to polysemy. In particular, REASON is annotated only on LUs inherently indicating volitional causation, either because a subcategorization frame making the agent of the action explicit exists, as with *The rain forced me to go*, or simply because the LU cannot describe a non-volitional CAUSATION, as is the case of the French connective *puisque* (since) (Degand and Pander Maat, 2003). In other words, a volitional cause expressed using a LU not implying volition per se, as in *I left because of the rain* is considered a case of CAUSATION, differently from the discourse annotation tradition (Degand and Pander Maat, 2003). A few LUs whose semantics are a priori volitional but without specific subcategorization, like the nouns *raison* (reason) and *motif* (motive) turned out to be ambiguous due to a few atypical non-volitional uses found in the corpus. Disambiguation is then easily guided by the nature of the effect, volitional or not.

Second, we faced a well-known difficulty with the annotation of ambiguous discourse connectives like *car* (because, for). Distinguishing objective from subjective causality, i.e., cases of CAUSATION, EVIDENCE, and even CAUSE.ENUNCIATION, is recognized as the most difficult task in annotating causal discourse relations (see, e.g. the recent (Scholman et al., 2016) which confirms earlier studies). The classical abduction examples of EVIDENCE, such as *He loves her because he came back* in which the objective causal relation between facts is presented as reversed from an epistemic point of view are relatively easy to spot

for this precise reason, but turn out to be much less frequent than cases in which the objective and the subjective readings are congruent. Other relatively easy cases are found when the support to the proposition argued for is a full argument with a logical structure and cannot be confused with an event or situation located in time, as in *You have to go now because if you don't, you will be late*. For the remainder, following suggestions found in the literature, the identification of an EVIDENCE case relies on the presence of marks of reported cognition (*Je pense* – I think), modality (tense, adverbs like *peut-être* – maybe, *certainement* – certainly), evaluation (*N est important* – N is important), opinion on a proposed action or attitude (*il faut V* – one has to V), or a combination of these in the proposition argued for. An example from the corpus would be:

- (6) Je pense que c'est là [une initiative extrêmement importante]_{proposition}, **car** [les opposants au régime de Mugabe ont manifestement besoin de protection]_{support}.
(*I think it is a very important initiative, as opponents to Mugabe's regime obviously need protection.*)

Finally, remaining ambiguous cases are annotated at the lowest level in the subjectivity scale (Degand and Pandera Maat, 2003), thus favoring a CAUSATION reading over an EVIDENCE reading and the latter over a CAUSE-ENUNCIATION reading.

Nouns and verbs did not require such fine-grained disambiguation rules to distinguish CAUSATION from EVIDENCE, as only specific lemmas like *attester*, *prouver* or *indiquer* belong to the lexicon of EVIDENCE, confirming earlier findings that generic causal verbs only denote objective causation (Stukker et al., 2008). In those cases though, disambiguation between the causal domain and the communication domain is required.

The total number of sentences premarked with a lexical unit potentially signalling a causality frame is about 7500 sentences. Of these, roughly 70% were judged to actually contain a causality-related frame instance. There is of course a lot of variance according to lexical items, and the natural distribution would be slightly different since we applied a threshold on frequent units.

Besides core role fillers, the annotation includes cases which would be ignored in FN, as “non-local” instantiation of the roles⁴. Whenever a role filler was clearly mentioned in the sentence though non-locally, the annotators were asked to mark it nonetheless, since it provides interesting semantic information (e.g. selectional preferences). For the same reason antecedents of anaphoric mentions were also annotated, if the antecedent is present in the same sentence as the frame, on top of the anaphor.

More details about the general annotation campaign and choices made can be found in (Djemaa et al., 2016).

Table 1 shows the inter-annotator agreement (IAA) with respect to the disambiguation of a frame-evoking item be-

tween its potential senses, and with respect to the annotation of role fillers. With respect to other domains within the project, the IAA for frames is slightly lower as the average on the whole project is 85.8% vs 78.4 for the causality domain (Djemaa et al., 2016), as is role annotation (77 vs 73). Besides the inherent difficulty of causality annotation, the domain shows the most variety in parts of speech for frame-evoking lexical items, while we observed a much better agreement on verbs and nouns than on adverbial markers (for frames), and a much better agreement for verbal roles than all other roles.

Table 2 shows the distribution of causal frames among the adjudicated annotations.

Nb of frames (annotator 1)	6982
Nb of frames (annotator 2)	6779
F-score (frames)	78.454
F-score (partial role)	80.260
F-score (exact role)	73.907

Table 1: Summary of the annotated data: number of annotated frame occurrences for lexical items with at least one causal sense, and inter-annotator agreement.

Frame	Occ.	%
Causation	1637	42,60
Evidence	701	18,24
Reason	427	11,11
Contingency+Objective influence*	419	10,90
Cause to start+Launch process*	264	6,87
Response	171	4,45
Make possible to do	110	2,86
Attributing cause*	44	1,14
Explaining the facts	40	1,04
Preventing	24	0,62
Cause enunciation*	6	0,16
Total	3843	100

Table 2: Break-up of actual causal frame occurrences in the adjudicated data; * indicates new or merged frames.

4. Conclusion

We present a modified set of FrameNet frames for the expression of causality, along with annotated frame instances on French treebanks. We will release the corresponding lexicon, frame descriptions and structure, and a disambiguation guide, along with the annotated French corpus, at the project website <http://alpage.inria.fr/asfalda>.

5. Acknowledgements

This work was funded by the French National Research Agency (ASFALDA project ANR-12-CORD-023), and supported by the French Investissements d’Avenir - Labex EFL program (ANR-10-LABX-0083). We gratefully acknowledge the work of our annotators, Anny Soubeille and Marjorie Raufast.

⁴More precisely, FN uses a fixed set of non-local instantiations that are to be annotated, namely raising, control, tough-movement and relativization (Ruppenhofer et al., 2006).

6. Bibliographical References

- Abeillé, A. and Barrier, N. (2004). Enriching a french treebank. In *Proc. of LREC'04*, Lisbon, Portugal.
- Atallah, C. (2014). *Analyse de relations de discours causales en corpus : étude empirique et caractérisation théorique*. Ph.D. thesis, Universit de Toulouse II.
- Baker, C. F., Fillmore, C. J., and Lowe, J. B. (1998). The berkeley framenet project. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics, Volume 1*, pages 86–90, Montreal, Quebec, Canada, August. Association for Computational Linguistics.
- Bethard, S. and Martin, J. H. (2008). Learning semantic links from a corpus of parallel temporal and causal relations. In *Proceedings of ACL-08: HLT, Short Papers*, pages 177–180, Columbus, Ohio, June. Association for Computational Linguistics.
- Burchardt, A., Erk, K., Frank, A., Kowalski, A., Padó, S., and Pinkal, M. (2006). Salto – a versatile multi-level annotation tool. In *Proceedings of LREC 2006*, Genoa, Italy.
- Candito, M. and Seddah, D. (2012). Le corpus sequoia : annotation syntaxique et exploitation pour l’adaptation d’analyseur par pont lexical. In *Proc. of TALN 2012 (in French)*, Grenoble, France, June.
- Candito, M., Amsili, P., Barque, L., Benamara, F., de Chalendar, G., Djemaa, M., Haas, P., Huyghe, R., Mathieu, Y. Y., Muller, P., Sagot, B., and Vieu, L. (2014). Developing a french framenet: Methodology and first results. In Nicoletta Calzolari, et al., editors, *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 1372–1379, Reykjavik, Iceland, May. European Language Resources Association (ELRA). ACL Anthology Identifier: L14-1411.
- Carlson, L., Marcu, D., and Okurowski, M. (2007). Rst discourse treebank ldc2002t07. Web download.
- Chklovski, T. and Pantel, P. (2004). Verbocean: Mining the web for fine-grained semantic verb relations. In Dekang Lin et al., editors, *Proceedings of EMNLP 2004*, pages 33–40, Barcelona, Spain, July. Association for Computational Linguistics.
- Davidson, D. (1969). The individuation of events. In N. Rescher, editor, *Essays in Honour of Carl Hempel*. Dordrecht: Reidel.
- Degand, L. and Pander Maat, H. (2003). A contrastive study of dutch and french causal connectives on the speaker involvement scale. In Arie Verhagen et al., editors, *Usage-based Approaches to Dutch*, pages 175–199. LOT, Utrecht.
- Djemaa, M., Candito, M., Muller, P., and Vieu, L. (2016). Corpus annotation within the French FrameNet: methodology and results. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*.
- Do, Q., Chan, Y. S., and Roth, D. (2011). Minimally supervised event causality identification. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 294–303, Edinburgh, Scotland, UK., July. Association for Computational Linguistics.
- Dunietz, J., Levin, L., and Carbonell, J. (2015). Annotating causal language using corpus lexicography of constructions. In *Proceedings of The 9th Linguistic Annotation Workshop*, pages 188–196, Denver, Colorado, USA, June. Association for Computational Linguistics.
- Gordon, A., Kozareva, Z., and Roemmele, M. (2012). Semeval-2012 task 7: Choice of plausible alternatives: An evaluation of commonsense causal reasoning. 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 394–398, Montréal, Canada, 7-8 June. Association for Computational Linguistics.
- 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*, pages 1172–1181, Singapore, August. Association for Computational Linguistics.
- Hovy, E. and Maier, E. (1992). Parsimonious or profligate: How many and which discourse structure relations? Technical Report RR-93-373, USC Information Sciences Institute.
- Khoo, C. S. G., Chan, S., and Niu, Y. (2000). Extracting causal knowledge from a medical database using graphical patterns. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*, Hong-Kong, October. Association for Computational Linguistics.
- Kozareva, Z. (2012). Cause-effect relation learning. In *Workshop Proceedings of TextGraphs-7: Graph-based Methods for Natural Language Processing*, pages 39–43, Jeju, Republic of Korea, July. Association for Computational Linguistics.
- Mak, W. and Sanders, T. (2013). The role of causality in discourse processing: effects on expectation and coherence relations. *Language and Cognitive processes*, 28(9):1414–1437.
- Mirza, P. and Tonelli, S. (2014). An analysis of causality between events and its relation to temporal information. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2097–2106, Dublin, Ireland, August. Dublin City University and Association for Computational Linguistics.
- Mouton, C., de Chalendar, G., and Richert, B. (2010). Framenet translation using bilingual dictionaries with evaluation on the English-French pair. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta, may. European Language Resources Association (ELRA).
- nan Cao, Y., Zhang, P., Guo, J., and Guo, L. (2014). Mining large-scale event knowledge from web text. *Procedia Computer Science*, 29:478 – 487. 2014 International

- Conference on Computational Science.
- Neeleman, A. and van de Koot, H. (2012). The linguistic expression of causation. In Paramita Mirza et al., editors, *The Theta System: Argument Structure at the Interface*, chapter 2, pages 20–51. Oxford University Press.
- Oh, J.-H., Torisawa, K., Hashimoto, C., Sano, M., De Saeger, S., and Ohtake, K. (2013). Why-question answering using intra- and inter-sentential causal relations. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1733–1743, Sofia, Bulgaria, August. Association for Computational Linguistics.
- Padó, S. (2007). *Cross-Lingual Annotation Projection Models for Role-Semantic Information*. Ph.D. thesis, Saarland University. MP.
- Prasad, R., Dinesh, N., Lee, A., Miltsakaki, E., Robaldo, L., Joshi, A., and Webber, B. L. (2008). The Penn Discourse TreeBank 2.0. In *Proceedings of LREC 2008*.
- Riaz, M. and Girju, R. (2013). Toward a better understanding of causality between verbal events: Extraction and analysis of the causal power of verb-verb associations. In *Proceedings of the SIGDIAL 2013 Conference*, pages 21–30, Metz, France, August. Association for Computational Linguistics.
- Riaz, M. and Girju, R. (2014). Recognizing causality in verb-noun pairs via noun and verb semantics. In *Proceedings of the EACL 2014 Workshop on Computational Approaches to Causality in Language (CatoCL)*, pages 48–57, Gothenburg, Sweden, April. Association for Computational Linguistics.
- Roze, C., Danlos, L., and Muller, P. (2012). LEXCONN: a French lexicon of discourse connectives. *Discours*, 10.
- Ruppenhofer, J., Ellsworth, M., Petruck, M. R., Johnson, C. R., and Scheffczyk, J. (2006). *FrameNet II: Extended Theory and Practice*. International Computer Science Institute, Berkeley, California. Distributed with the FrameNet data.
- Sanders, T. and Stukker, N. (2012). Causal connectives in discourse: A cross-linguistic perspective – introduction to the special issue. *Journal of Pragmatics*, 44(2).
- Scholman, M., Evers-Vermeul, J., and Sanders, T. (2016). A step-wise approach to discourse annotation: Towards a reliable categorization of coherence relations. *Dialogue and Discourse*, 7(2):1–28.
- Spooren, W., Sanders, T., Huiskes, M., and Degand, L. (2010). Subjectivity and causality: A corpus study of spoken language. In S. Rice et al., editors, *Empirical and Experimental Methods in Cognitive/Functional Research*, pages 241–255. CSLI Publications.
- Stukker, N., Sanders, T., and Verhagen, A. (2008). Causality in verbs and in discourse connectives: Converging evidence of cross-level parallels in dutch linguistic categorization. *Journal of Pragmatics*, 40:1296–1322.
- Sweetser, E. (1990). *From etymology to pragmatics: Metaphorical and cultural aspects of semantic structure*. Cambridge University Press, Cambridge.
- Szpektor, I. and Dagan, I. (2008). Learning entailment rules for unary templates. In *Proceedings of the 22nd International Conference on Computational Linguistics (Coling 2008)*, pages 849–856, Manchester, UK, August. Coling 2008 Organizing Committee.
- Wolff, P. and Song, G. (2003). Models of causation and the semantics of causal verbs. *Cognitive Psychology*, 47(3):276–332.
- Zufferey, S. and Cartoni, B. (2012). English and french causal connectives in contrast. *Languages in contrast*, 12(2):232–250.