



A Repository of Frame Instance Lexicalizations for Generation

Valerio Basile

► **To cite this version:**

Valerio Basile. A Repository of Frame Instance Lexicalizations for Generation. Proceedings of the 2nd International Workshop on Natural Language Generation and the Semantic Web, Sep 2016, Edinburgh, United Kingdom. <hal-01389398>

HAL Id: hal-01389398

<https://hal.inria.fr/hal-01389398>

Submitted on 28 Oct 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Repository of Frame Instance Lexicalizations for Generation

Valerio Basile

Université Côte d’Azur, Inria, CNRS, I3S, France

valerio.basile@inria.fr

Abstract

Robust, statistical Natural Language Generation from Web knowledge bases is hindered by the lack of text-aligned resources. We aim to fill this gap by presenting a method for extracting knowledge from natural language text, and encode it in a format based on frame semantics and ready to be distributed in the Linked Open Data space. We run an implementation of such methodology on a collection of short documents and build a repository of frame instances equipped with fine-grained lexicalizations. Finally, we conduct a pilot study to investigate the feasibility of an approach to NLG based on said resource. We perform error analysis to assess the quality of the resource and manually evaluate the output of the NLG prototype.

1 Introduction

Statistical Natural Language Generation, generally speaking, is based on learning a mapping between natural language expressions (words, phrases, sentences) and abstract representations of their meaning or syntactic structure. In fact, such representations vary greatly in their degree of abstraction, from shallow syntactic trees to full-fledged logical formulas, depending on factors like downstream applications and the role of the generation module in a larger framework.

In order to be useful for statistical generation, the abstract representation needs to be aligned with the surface form. Depending on the format, the level of abstraction and the target degree of granularity of the alignment, it may be more or less straightforward to produce a collection of pairs <abstract representation, surface form>. Moreover, statistical methods typically

need a large number of examples to properly learn a mapping and generalize efficiently.

While several resources have been successfully employed as training material for statistical NLG (see the related work section), they lack a direct link with world knowledge. Linked Open Data resources, in particular general knowledge bases such as DBpedia¹, on the other hand, are not straightforward to use as a basis for generation, while at the same time they are rich in extralinguistic information such as type hierarchy and semantic relations. Having the entities and concepts of an abstract meaning representation linked to a knowledge base allows a generator to use all the information coming from links to other resources in the LOD cloud. Such kind of input to a NLG pipeline is therefore richer than word-based structures, although its increased level of abstraction makes the generation process more complex.

Shifting the level of abstraction, the representation format must be changed accordingly. In the case of many formats proposed in the literature (e.g., the format of the Surface Realization shared task), the input for NLG is made of structures closely resembling sentences. The notion of sentence, however, might not be adequate anymore when the abstract representation of meaning aims to be fit for the standards of the Web. A good compromise is a representation based on frame semantics (Fillmore, 1982). A *frame* is a unit of meaning denoting a situation of a particular type, e.g., *Operate_vehicle*. Attached to the frame there are a number of *frame elements*, indicating roles that the entities involved in the frame can play, e.g., *Driver* or *Vehicle*. Rouces et al. (2015) proposes a LOD version of frame semantics implemented in the resource called FrameBase, essentially a scheme for representing instances of frames and frame ele-

¹<http://dbpedia.org>

ments in a Web-based format. The FrameBase project also produced a repository of instances created by automatically translating existing Web resources. Moreover, they made available a large set of (*de*)reification rules, that is, bidirectional rules to convert between binary relations and frame-based representations. For instance, the binary relation *drivesVehicle* can be transformed by a reification rule into a *Operate_vehicle* frame with the two members of the original relation filling in the roles of *Driver* and *Vehicle*. The reification mechanism provides an interesting use case for NLG: if a system is able to generate natural language from a frame instance, then it is also able to generate from the corresponding binary relation.

In this paper, we present an ongoing work towards the construction of a domain-agnostic, LOD-compliant knowledge base of semantic frame instances. Frames, roles and entities are aligned to natural language words and phrases that express them, extracted from a large corpus of text. Thanks to this alignment, the resource can be used to create lexicalizations for new, unseen configurations of entities and frames.

2 Related Work

Several resources exist that have been used to train a statistical generator to learn lexicalizations for various types of representations. The Surface Realization Shared Task (Belz et al., 2011), for instance, provides a double dataset of shallow and deep input representations obtained by preprocessing the CoNLL 2008 Shared Task data (Surdeanu et al., 2008). Resources used for NLG include including the Penn Treebank (Marcus et al., 1993) for Probabilistic Lexical Functional Grammar (Cahill and Genabith, 2006) and CCGBank (Hockenmaier and Steedman, 2007) for Combinatory Categorical Grammar syntax trees (White et al., 2007). More recently, the Groningen Meaning Bank (Basile et al., 2012) has been proposed as a resource for NLG from abstract meaning representations, leveraging the fine-grained alignment between logical forms and their respective surface forms given by the Discourse Representation Graph formalism (Basile and Bos, 2013).

The process of generating natural language from databases of structured information, including ones following Web standards, has been studied in the past, although often in specific application-oriented contexts. Bouayad-Agha et

al. (2012) propose an architecture as a basis for generation made of three RDF/OWL ontologies, separating the domain knowledge from the communication knowledge. Gyawali and Gardent (2014) propose a statistical approach to NLG from knowledge bases based on tree adjoining grammars. WordNet is relatively less used for generation purposes. Examples of the use of WordNet in the context of NLG include the methods to address specific NLG-related tasks proposed by Jing (1998) and the algorithm for lexical choice of Basile (2014).

3 Aligning Text and Semantics

Basile and Bos (2013) devise a strategy to align arbitrary natural language expressions to formal representation of their meaning, encoded as Discourse Representation Structures (DRS, Kamp and Reyle (1993)). DRSs are logical formulas comprising predicates and relations over discourse referents. For the English language, we are able to obtain DRSs for a given text using the C&C tools collection of linguistic analysis tools (Curran et al., 2007), which includes Boxer (Bos, 2008), a rule-based system that builds DRSs on top of the CCG parse tree produced by the C&C parser. Boxer implements *Neo-davidsonian* representations of meaning, that is, formulas centered around *events* to which participant entities are connected by filling *thematic roles*. Figure 1 shows an example of DRS for the sentence “A robot is driving the car” as produced by Boxer. In this example the Neo-davidsonian semantics is evident: the ROBOT is the AGENT of the event DRIVE, while the CAR is the THEME.

e_1	x_1	x_2
ROBOT	(x_1)	
DRIVE	(e_1)	
CAR	(x_2)	
AGENT	(e_1, x_1)	
THEME	(e_1, x_2)	

Figure 1: DRS representing the meaning of the sentence “A robot is driving the car”

The alignment method proposed by Basile and Bos (2013) is based on a translation of format from DRS into a Discourse Representation Graph (DRG), where the semantic information is preserved but expressed in a flat, non recursive formalism. The surface form is then aligned at the

word level to the appropriate tuples. Figure 2 shows the DRG corresponding to the DRS in Figure 1, where the alignment with the surface form is contained in the two rightmost columns. For the details of how the alignment is encoded we refer the reader to the aforementioned paper (Basile and Bos, 2013).

k_1	referent	x_1	1	[A]
k_1	referent	e_1		
k_1	referent	x_2	1	[the]
k_1	event	DRIVE		
k_1	concept	ROBOT		
k_1	role	AGENT		
k_1	concept	CUSTOMER		
k_1	role	THEME		
ROBOT	instance	x_1	2	[robot]
DRIVE	instance	e_1	2	[is, driving]
AGENT	internal	e_1	1	
AGENT	external	x_1		
CAR	instance	x_2	2	[car]
THEME	internal	e_1	3	
THEME	external	x_2		

Figure 2: DRG aligned with the surface form, representing the meaning of the sentence “A robot is driving the car”.

In order for the semantic representations, and their alignment to the surface, to be useful in contexts such as knowledge representation and automatic reasoning, these logical forms need to be linked to some kind of knowledge base. Otherwise, the predicate symbols in a DRG like the one depicted in Figure 2 are just interchangeable symbols (although Boxer uses lemmas for predicate names) devoid of meaning.

Popular resources in the LOD ecosystem are well-suited for serving as knowledge bases for grounding the symbols: WordNet (Miller, 1995) can be used to represent concepts and events, while DBpedia has a very large coverage for named entities. FrameNet (Baker et al., 1998), an inventory of frames and frame elements inspired by Fillmore’s frame semantics (Fillmore, 1982), has a structure that superimposes easily to the neo-Davidsonian semantics of Boxer’s DRGs. The inventory of thematic roles used by Boxer is taken from VerbNet (Schuler, 2005). By linking the discourse referents representing concepts in a DRG to WordNet synsets, entities to DBpedia and events to FrameNet frames we are able to extract complete representations of frames from natural language text linked to LOD knowledge bases.

4 Collecting Frame Lexicalizations

We developed a pipeline of NLP tools to automatically extract instances of frames from the text. The pipeline comprises the C&C tools and Boxer, a module for word sense disambiguation and a module for entity linking. The two latter modules can be configured to use different external software to perform their task.

The analysis of a text consists in the following steps:

1. Run the C&C tools and Boxer, saving both its XML and DRG output. The XML output of Boxer contains, for each predicates of the DRS that has been constructed, a link to the part of the surface form that introduced it.
2. Run the WSD and entity linking components, preserving the same tokenization. The software then uses the links to the text provided by Boxer to map the word senses and DBpedia entities to the DRS predicates.
3. The word senses corresponding to events are mapped to FrameNet frames, using the mapping provided by Rouces et al. (2015). The VerbNet roles are converted into FrameNet roles using the mapping provided by Loper et al. (2007).
4. The partial surface forms in the DRG output of Boxer are attached to the frames, semantic roles and frame elements.




Figure 3: Architectural Scheme of KNEWS.

This pipeline is implemented in the KNEWS system, available for download at <https://github.com/valeriobasile/learningbyreading>. In the following paragraphs we describe the internal details of the components of KNEWS.

Semantic parsing The semantic parsing module employs the C&C tools and Boxer to process the

input text and output a complete formal representation of its meaning. The C&C pipeline of statistical NLP tools includes a tokenizer, a lemmatizer, named entity and part-of-speech tagger, and a parser that creates a Combinatorial Categorical Grammar representation of the natural language syntax. Boxer builds a DRS on top of the CCG analysis. The predicates of a DRS are expressed over a set of *discourse referents* representing entities, concepts and events. Such structures contain, among other information, predicates representing the roles of the entities with respect to the detected events, e.g., *event(A)*, *entity(B)*, *agent(A,B)* to represent B playing the role of the *agent* of the event A.

Word sense disambiguation and Entity Linking

KNEWS uses WordNet to represent concepts and events, DBpedia for named entities, and FrameNet’s frames to represent events, integrating the mapping with the WordNet synsets provided by FrameBase. The inventory of thematic roles used by Boxer is taken from VerbNet (Schuler, 2005), while KNEWS employs the mapping provided by SemLinks (Palmer, 2009) to link them (whenever possible) to FrameNet roles. KNEWS can be configured to use either UKB (Agirre and Soroa, 2009) or Babelfy (Moro et al., 2014) to perform the word sense disambiguation, and DBpedia Spotlight (Daiber et al., 2013) or Babelfy for entity linking.

Output modes KNEWS’s default output consists of frame instances, sets of RDF triples that contain a unique identifier, the type of the frame, the thematic roles involved in the instance, and the concepts or entities that fill the roles. The format follows the scheme of FrameBase, which offers the advantage of interoperability with other resources in the Linked Open Data cloud, as well as the possibility of using FrameBase’s (de)reification rules to automatically generate a large number of binary predicates. An example of frame instance, extracted from the sentence “A robot is driving the car.” is given in Figure 4.

This output mode of KNEWS has been employed in Basile et al. (2016) to create a repository of general knowledge about objects.

For the purpose of NLG, we extended KNEWS with a new output mode, similar to the previous one (frame instances) with the difference that it contains as additional information the alignment

with the text. We exploit the DRG output of Boxer to link the discourse referents to surface forms, i.e., span of the original input text, resulting in the word-aligned representation shown in Figure 5. This new output mode of KNEWS consist of an XML list of *frameinstance* elements. Each frame instance is equipped with its complete lexicalization (the *instancelexicalization* tag), the incomplete surface form associated with the event (the *framelexicalization* tag) and a sequence of *frameelements*. A *frameelement* represent a role in the frame instance. The *concept* tag contains a DBpedia or Wordnet resource (depending on the output of the disambiguation module), a lexicalization of the role filler (the *conceptlexicalization* tag), and the incomplete surface form obtained by composing the surface forms of the role filler and the frame. In the next section we describe an automatically built resource created by parsing text with this configuration of KNEWS.

KNEWS has also an additional output mode: First-order Logic. With this output mode, KNEWS is able to generate first-order logic formulae representing the natural language text given as input. The symbols for the predicates are Wordnet symbols, allowing the output of KNEWS to be integrated with a reasoning engine, e.g., to select background knowledge in a much more focused manner, as proposed in Furbach and Schon (2016).

5 Evaluation

In order to test our approach to knowledge extraction, we parsed a corpus of short texts, taken from the ESL Yes website of material for English learners.² We find this data particularly apt in the more general context of extracting general knowledge from text, being made of short, clear sentences about simple and generic topics. The corpus comprises 725 short stories, that we divided into 14,140 sentences. Parsing the ESL Yes corpus with KNEWS we collected 30,217 frame instances (420 unique frames), 1,455 concepts (1,201 WordNet synsets and 254 DBpedia entities) filling in 41,945 roles (161 unique roles). 29,409 role instances could not be mapped to FrameNet, so they are expressed by one of 18 VerbNet roles.

We evaluate the information extraction methodology by assessing the quality of this automatically produced resource. For each frame instance,

²<http://www.eslyes.com/>

```

@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix fb: <http://framebase.org/ns/> .
@prefix dbr: <http://dbpedia.org/resource/> .
@prefix wn: <http://wordnet-rdf.princeton.edu/wn31/> .

fb:fi-Operate_vehicle_dc59afa6 rdf:type fb:frame-Operate_vehicle-drive.v .
fb:fi-Operate_vehicle_dc59afa6 fb:fe-Driver dbr:Robot .
fb:fi-Operate_vehicle_dc59afa6 fb:fe-Vehicle wn:02961779-n .

```

Figure 4: RDF triples extracted by KNEWS from the sentence “A robot is driving the car”, constituting one frame instance.

```

<frameinstance id="Operate_vehicle_dc59afa6"
  type="Operate_vehicle-drive.v" internalvariable="e1">
  <framelexicalization>k3:x1 is driving k3:x2</framelexicalization>
  <instancelexicalization>The robot is driving the car .</instancelexicalization>
  <frameelements>
    <frameelement role="Driver" internalvariable="x1">
      <concept>http://dbpedia.org/resource/Robot</concept>
      <rolelexicalization>The robot is driving x2</rolelexicalization>
      <conceptlexicalization>The robot</conceptlexicalization>
    </frameelement>
    <frameelement role="Vehicle" internalvariable="x2">
      <concept>http://wordnet-rdf.princeton.edu/wn31/02961779-n</concept>
      <rolelexicalization>x1 is driving the car .</rolelexicalization>
      <conceptlexicalization>the car .</conceptlexicalization>
    </frameelement>
  </frameelements>
</frameinstance>

```

Figure 5: XML output of KNEWS describing a frame instance extracted from the sentence “A robot is driving the car”.

if all the information is present and complete, it should be possible to recreate the instance lexicalization by applying the composition method of Basile and Bos (2013). The incomplete surface forms corresponding to the frame and the frame elements are automatically composed and compared to the original frame lexicalization. We ran this evaluation procedure on the resource and found 7,366 instances are correctly regenerated, that is, about one in four instances. Of the remaining instances, 11,996 present incorrect instance lexicalizations, usually containing variables instead of being complete surface forms. These occurrences are caused by misalignments in the representation produced by Boxer, so that the composition algorithm cannot recreate the original surface form. For instance, for the sentence “The mother gave her baby a red apple”, the lexicalized DRG produced by Boxer, when the composition algorithm is applied to it, produces “The mother gave k5:x3 baby k4:x2”. We also found that in 5,211 cases the presence subordination prevents the realization algorithm from working correctly, because no lexicalization is found for the discourse referent corresponding to the subordinate clause. In 1,865 cases, issues are caused by the presence of phrasal verbs (e.g. “He picked up his clothes”) or adverbs, which are analyzed by Boxer using the

Table 1: Error analysis of the automatically produced, text-aligned frame instance collection, broken down by number of frame elements.

roles	1	2	3	all
correct	4,774	2,374	218	7,366
subordination	4,824	368	19	5,211
adverb	1,288	561	16	1,865
realization	5,885	5,508	603	11,996
other	2,672	1,009	98	3,779
total	19,443	9,820	954	30,217

relation *manner* between the event and the adverb or proposition, thus like in the previous case no lexicalization is found for all the discourse referents. Finally, 3,779 instances failed the test due to a variety of reasons, e.g., failure of the entity linking module or wrong syntactic analysis. Table 1 summarizes the findings exposed so far, also broken down by the number of frame elements in the frame instances.

When increasing number of frame elements per frame instance, the issues with subordinate constructions dramatically decreases: they amount to 24% of the cases with one frame elements, 3% and 1% with two and three frame elements respectively. Conversely, wrong realizations due to representation misalignments tends to get worse, involving from 30% of the instances with one frame elements to 56% with two, to 39% with three.

6 Generation of Frame Lexicalizations

The first and most obvious use for the resource presented here in the context of NLG is given by the set of lexicalizations it provides for concepts and entities. In the example in Figure 5, for instance, the DBpedia entity `Robot` is lexicalized as “A robot” and the synset `02961779-n` as “the car”. Moreover, the frame is also given the lexicalization with two open variables “ x_1 is driving x_2 ”. Indeed, the surface forms provided by the DRG can be *incomplete*, that is, containing variables that can be used to compose a full surface form from the single ones corresponding to the discourse referents, e.g., x_1 :“A robot” and e_1 :“ x_1 is driving x_2 ” compose to form e_1 :“A robot is driving x_2 ”, and so on.

This composition mechanism gives us the opportunity to devise a simple method to produce new frame lexicalizations. Given new concepts or entities with the respective lexicalizations and roles (e.g., Driver: “Valentino Rossi”, Vehicule: “the motorbike”), they can be replaced in the appropriate frame instance so that the variables x_1 and x_2 are linked respectively to “Valentino Rossi” and “the motorbike”. A subsequent step of composition will then yield the new frame lexicalization “Valentino Rossi is driving the motorbike”.

We developed a simple prototype in order to test this approach to NLG from frame instances. This prototype is based on the resource described in Section 5, restricted to the instances with exactly two frame elements and associated with a complete surface form. The procedure we use to evaluate the system is the following:

1. For each frame instance, produce four new frame instances by replacing one or both frame elements, either with similar concepts or with randomly chosen concepts.
2. Generate the lexicalization of the new frames by composing the frame lexicalization structure with the new concept lexicalizations.
3. For each of the four scenarios, select randomly one hundred instance lexicalization for the evaluation.
4. Manually inspect the selected lexicalizations according to three possible classes of fluency: *nonsensical* (the sentence is not grammatical and it does not make sense), *informative*

Table 2: Result of the manual evaluation of the NLG prototype based on the collection of lexicalized frame instances.

Replaced frame elements	Judgment nonsensical/informative/fluent
1, most similar	23/33/44
2, most similar	24/53/23
1, random	23/35/42
2, random	54/23/23

(the grammar contains mistakes but the information is clearly transmitted), and *fluent* (the lexicalization correctly conveys the input knowledge).

When we replace one frame element or both of them with similar concepts, we rely on the WUP similarity defined by Wu and Palmer (1994) for pairs of WordNet synsets, a measure of path distance weighted according to the depth of the WordNet taxonomy. We compute the WUP similarity for each pair of concepts in our collection and replace one or both frame elements with their most similar concepts. For example, the frame elements corresponding to the *Vehicle* in the frame instance in Figure 5 is associated with the concept <http://wordnet-rdf.princeton.edu/wn31/02961779-n> (car, automobile). This concept could be replaced, for the sake of the evaluation, by the similar concept (according to the WUP metric) <http://wordnet-rdf.princeton.edu/wn31/104497386-n> (truck), if this is also in the collection. A new lexicalization is then produced by composition “A robot is driving **the truck**”. The lexicalization for the replaced concepts is chosen as the most frequent lexicalization of that particular concept, to minimize the occurrence of awkward realizations like “A robot is driving **of the truck**”.

Note that we only judge fluency. An evaluation of adequacy or other content-oriented metrics should also take into account the input and would be more difficult to evaluate in this setting, since here the input is artificially produced by replacing elements of the frame instances.

The manual inspection of the produced frame instance lexicalizations resulted in the figures shown in Table 2. As expected, replacing both frame elements instead of just one leads to more errors in the realizations. This problem can be mitigated by increasing the coverage of the resource. With a larger collection, the chance of retrieving a

frame instance with at least one frame element in common with the new input is higher, thus there will be more cases where only one frame element is new. Interestingly, the choice of concepts to generate with respect to the frame (similar vs. random) does not seem to influence the outcome. The result of this pilot study are encouraging in that a sufficiently large number of correct realizations are produced by a simple mechanism. However, a more thorough evaluation is needed, especially with respect to the coverage (and thus the scalability) of our approach.

7 Conclusion and Future Work

In this paper we introduced a novel methodology to extract knowledge from text and encode it in formal structures compatible with the standards of the Web. Such structures are essentially instances of frames with their frame elements linked to concepts in Wordnet or DBpedia. This methodology is implemented in the freely available software package KNEWS. Next, we presented a collection of frame instances aligned with natural language, automatically created by parsing text for English learners. Finally, we propose a pilot study on how to use this resource to generate natural language from new frame instances.

In terms of future direction for this work, the low hanging fruit is the enlargement of the resource, which will lead to a higher number of “good” instances to use for direct generation (as shown in Section 6) and more data to use for a statistical approach to generation. Since the resource is produced automatically by parsing raw text with KNEWS, and natural language is abundant on the Web, this is a direction we intend to take in the foreseeable future.

The approach to NLG based on the collection of lexicalized frame instances introduced in NLG is at the preliminary work stage, and many refinements can be made to the algorithm. Given a new frame instance to generate, its frame elements could be matched to the lexicalization in the resource with more sophisticated methods, e.g., using distributional similarity.

As a possible extension to the resource, information such as lemma and number could be included in the lexicalization of concepts. With such information in place, the NLG algorithm could be interfaced with the SimpleNLG surface realization library (Gatt and Reiter, 2009) to produce more

fluent lexicalizations.

The main selling point of a large knowledge base aligned with text is that its size allows researchers to develop statistical methods to learn a mapping between the formally encoded knowledge and natural language. While this could be a very challenging enterprise, as highlighted by the work presented in Basile (2015), this work constitutes a first step in this direction.

References

- Eneko Agirre and Aitor Soroa. 2009. Personalizing pagerank for word sense disambiguation. In *Proceedings of the 12th Conference of the European Chapter of the Association for Computational Linguistics*, EACL '09, pages 33–41, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The berkeley framenet project. In *Proceedings of the 17th International Conference on Computational Linguistics - Volume 1*, COLING '98, pages 86–90, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Valerio Basile and Johan Bos. 2013. Aligning Formal Meaning Representations with Surface Strings for Wide-coverage Text Generation. In *ENLG 2013*, page 1.
- Valerio Basile, Johan Bos, Kilian Evang, and Noortje Venhuizen. 2012. Developing a large semantically annotated corpus. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC 2012)*, pages 3196–3200, Istanbul, Turkey.
- Valerio Basile, Elena Cabrio, and Fabien Gandon. 2016. Building a general knowledge base of physical objects for robots. In *The Semantic Web. Latest Advances and New Domains*.
- Valerio Basile. 2014. A lesk-inspired unsupervised algorithm for lexical choice from wordnet synsets. *The First Italian Conference on Computational Linguistics CLiC-it 2014*, page 48.
- Valerio Basile. 2015. *From logic to language: Natural language generation from logical forms*. Ph.D. thesis.
- Anja Belz, Michael White, Dominic Espinosa, Eric Kow, Deirdre Hogan, and Amanda Stent. 2011. The first surface realisation shared task: Overview and evaluation results. In *Proceedings of the 13th European Workshop on Natural Language Generation*, ENLG '11, pages 217–226, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Johan Bos. 2008. Wide-coverage semantic analysis with boxer. In *Semantics in Text Processing. STEP 2008 Conference Proceedings*, volume 1, pages 277–286.

- Nadjet Bouayad-Agha, Gerard Casamayor, Simon Mille, Marco Rospocher, Horacio Saggion, Luciano Serafini, and Leo Wanner. 2012. From ontology to nl: Generation of multilingual user-oriented environmental reports. In Gosse Bouma, Ashwin Ittoo, Elisabeth Métais, and Hans Wortmann, editors, *Natural Language Processing and Information Systems*, volume 7337 of *Lecture Notes in Computer Science*, pages 216–221. Springer Berlin Heidelberg.
- Aoife Cahill and Josef Van Genabith. 2006. Robust pcfg-based generation using automatically acquired lfg approximations. In *In Proceedings of the 44th ACL*.
- James R. Curran, Stephen Clark, and Johan Bos. 2007. Linguistically motivated large-scale nlp with c&c and boxer. In *Proceedings of the 45th Annual Meeting of the ACL on Interactive Poster and Demonstration Sessions, ACL '07*, pages 33–36, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Joachim Daiber, Max Jakob, Chris Hokamp, and Pablo N. Mendes. 2013. Improving efficiency and accuracy in multilingual entity extraction. In *Proceedings of the 9th International Conference on Semantic Systems (I-Semantics)*.
- Charles Fillmore. 1982. Frame semantics. *Linguistics in the morning calm*, pages 111–137.
- Ulrich Furbach and Claudia Schon. 2016. Commonsense reasoning meets theorem proving. In *Proceedings of Workshop on Bridging the Gap between Human and Automated Reasoning*.
- Albert Gatt and Ehud Reiter. 2009. Simplenlg: A realisation engine for practical applications. In *Proceedings of the 12th European Workshop on Natural Language Generation, ENLG '09*, pages 90–93, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Bikash Gyawali and Claire Gardent. 2014. Surface realisation from knowledge-bases. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 424–434, Baltimore, Maryland, June. Association for Computational Linguistics.
- Julia Hockenmaier and Mark Steedman. 2007. Ccg-bank: A corpus of ccg derivations and dependency structures extracted from the penn treebank. *Comput. Linguist.*, 33(3):355–396, September.
- Hongyan Jing. 1998. Usage of wordnet in natural language generation. In *Proceedings of the Joint 17th International Conference on Computational Linguistics 36th Annual Meeting of the Association for Computational Linguistics (COLING-ACL'98) workshop on Usage of WordNet in Natural Language Processing Systems*, pages 128–134.
- Hans Kamp and Uwe Reyle. 1993. *From Discourse to Logic. Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Kluwer, Dordrecht.
- Edward Loper, Szu ting Yi, and Martha Palmer. 2007. Combining lexical resources: Mapping between propbank and verbnet. In *In Proceedings of the 7th International Workshop on Computational Linguistics*.
- Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of english: The penn treebank. *COMPUTATIONAL LINGUISTICS*, 19(2):313–330.
- George A. Miller. 1995. Wordnet: A lexical database for english. *Commun. ACM*, 38(11):39–41, November.
- Andrea Moro, Alessandro Raganato, and Roberto Navigli. 2014. Entity Linking meets Word Sense Disambiguation: a Unified Approach. *Transactions of the Association for Computational Linguistics (TACL)*, 2:231–244.
- Martha Palmer. 2009. SemLink: Linking PropBank, VerbNet and FrameNet. In *Proceedings of the Generative Lexicon Conference*, Pisa, Italy, Sept.
- Jacobo Rouces, Gerard de Melo, and Katja Hose. 2015. Framebase: Representing n-ary relations using semantic frames. In Fabien Gandon, Marta Sabou, Harald Sack, Claudia d’Amato, Philippe Cudr-Mauroux, and Antoine Zimmermann, editors, *ESWC*, volume 9088 of *Lecture Notes in Computer Science*, pages 505–521. Springer.
- Karin Kipper Schuler. 2005. *Verbnet: A Broad-coverage, Comprehensive Verb Lexicon*. Ph.D. thesis, Philadelphia, PA, USA. AAI3179808.
- Mihai Surdeanu, Richard Johansson, Adam Meyers, Lluís Màrquez, and Joakim Nivre. 2008. The conll-2008 shared task on joint parsing of syntactic and semantic dependencies. In *Proceedings of the Twelfth Conference on Computational Natural Language Learning, CoNLL '08*, pages 159–177, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Michael White, Rajakrishnan Rajkumar, and Scott Martin. 2007. Towards broad coverage surface realization with ccg. In *In Proc. of the Workshop on Using Corpora for NLG: Language Generation and Machine Translation (UCNLG+MT)*.
- Zhibiao Wu and Martha Palmer. 1994. Verbs semantics and lexical selection. In *Proceedings of the 32Nd Annual Meeting on Association for Computational Linguistics, ACL '94*, pages 133–138, Stroudsburg, PA, USA. Association for Computational Linguistics.