Proceedings of the International FrameNet Workshop 2020: Towards a Global, Multilingual FrameNet, pages 31–40
Language Resources and Evaluation Conference (LREC 2020), Marseille, 11–16 May 2020

Combining Conceptual and Referential Annotation to Study Variation in Framing

Marten Postma, Levi Remijnse, Filip Ilievski, Antske Fokkens, Sam Titarsolej, Piek Vossen
Vrije Universiteit Amsterdam
De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands
{m.c.postma,l.remijnse,antske.fokkens,piek.vossen}@vu.nl
ilievski@isi.edu, s.titarsolej@gmail.com

Abstract

We introduce an annotation tool whose purpose is to gain insights into variation of framing by combining FrameNet annotation with referential annotation. English FrameNet enables researchers to study variation in framing at the conceptual level as well through its packaging in language. We enrich FrameNet annotations in two ways. First, we introduce the referential aspect. Secondly, we annotate on complete texts to encode connections between mentions. As a result, we can analyze the variation of framing for one particular event across multiple mentions and (cross-lingual) documents. We can examine how an event is framed over time and how core frame elements are expressed throughout a complete text. The data model starts with a representation of an event type. Each event type has many incidents linked to it, and each incident has several reference texts describing it as well as structured data about the incident. The user can apply two types of annotations: 1) mappings from expressions to frames and frame elements, 2) reference relations from mentions to events and participants of the structured data.

Keywords: FrameNet, reference, annotation tool

1. Introduction

We construct narratives to describe events in the world around us. The language that we use in those narratives forms a lens that filters the actual components of those events, e.g., their time, location, and participants, according to our perspectives. This way, narratives function not only as structured collections of informative references to the event, but also as collections of conceptual representations of that same event. For instance, texts describing the attack on the World Trade Center express their references in various linguistic forms: '9/11', 'September 11 attacks', 'the 2001 attacks' (all of which are timestamps with different specificity), 'a series of four coordinated terrorist attacks' (focus on the organizational aspect), 'destruction of the towers in America' (focus on the damaging aspect), etc. This set of references is a small share of all the various references in a growing portion of written texts. With multiple texts written in different languages about a single real-world event, one could analyze variation of framing of an event by combining the conceptual and referential information. To perform such an analysis, we need both semantic resources to describe this conceptual information, and information about the components of the real-world event.

English FrameNet (Ruppenhofer et al., 2006) brought conceptual framing research to a computational setting. The English lexicon made it possible to gain insight into the relationship between lexical items and the semantic frames that they evoke. English FrameNet has also motivated researchers to create FrameNets in other languages such as in Japanese (Ohara et al., 2004), German (Burchardt et al., 2009), Swedish (Heppin and Gronostaj, 2012), Brazilian Portuguese (Laviola et al., 2017), Spanish (Subirats and Sato, 2003), French (Djemaa et al., 2016), Hebrew (Hayoun and Elhadad, 2016), and Latvian (Gruzitis et al., 2018). Multiple annotation efforts resulted in many corpora and also served as training, development, and test data to train FrameNet-based Semantic Role Labelers.

The majority of the described efforts have mainly investigated frame annotations at the sentence level, as already observed by Fillmore, evidenced by the following quote: "since FrameNet has been working mainly on single sentences and has done nothing (yet) on connections within whole texts, the FrameNet database has nothing direct to offer." (Andor, 2010, p.168)

We aim at combining FrameNet annotations with referential annotations in order to analyze framing variation in texts describing an event. For this we need to extend FrameNet annotations to the discourse level. Following the data-to-text method described in Vossen et al. (2018), we make use of the data acquisition platform described in Vossen et al. (2020) to enable this type of research, for which we require:

1. a referential representation of an incident, i.e., an event instance such as *the 2012 Slovenian presidential election*, with structured information about the location, time, and participants of the incident. 2. each incident to be tagged with one or more event types, e.g., *election*. This makes it possible to generalize over incidents of the same type to learn which frames are typical. 3. different texts that make reference to the same incident, possibly written in multiple languages with varying document creation times and from different sources, which provides us with insights into cross-lingual differences, source perspectives and the impact of historical distance to the incident time (Cybulska and Vossen, 2011). 4. an environment for efficient and consistent FrameNet and reference annotation to (given) structured data. This makes it possible to consider the framing of the incident throughout all texts that make reference to it as a discourse unit.

In this paper, we introduce an annotation tool in which both structured data about an incident and many reference texts describing that one incident are simultaneously presented to the user. This interface enables both FrameNet-based anno-

tations as well as referential linking to the incident that the reference texts make reference to. The analysis of conceptual and referential framing enriches research into variation in framing beyond the level of sentences and across different types of reference texts and languages.
This paper is structured as follows. In Section 2., we introduce English FrameNet and the related work on frame annotation, followed by a discussion on combining conceptual and referential annotation in Section 3. We introduce the annotation tool in Section 4. Finally, we discuss the possibilities of the tool and future plans in Section 5., and conclude the paper in Section 6.

2. Background

This section introduces the theoretical notions and implementations of FrameNet. Subsection 2.1. describes the relevant terminology and basic principles of frame semantics. In Subsection 2.2., we provide a brief overview of currently available frame annotation tools.

2.1. FrameNet

Frame semantics is a theory of linguistic meaning that assumes that the meaning of words is (partially) activated through the *frames* that they evoke (Ruppenhofer et al., 2006). A frame is a schematic representation of a concept, which is triggered by a lexical unit. This lexical unit is the sense of an expression in spoken or written discourse. For the purpose of this paper, we model these conceptual relationships using RDF, as displayed in Figure 1. In this figure, the expression 'kidnapped' is disambiguated to a lexical unit via the *ontolex:sense* relationship (McCrae et al., 2017). The lexical unit evokes the frame **Kidnapping** via the *ontolex:evokes* relationship.[1]

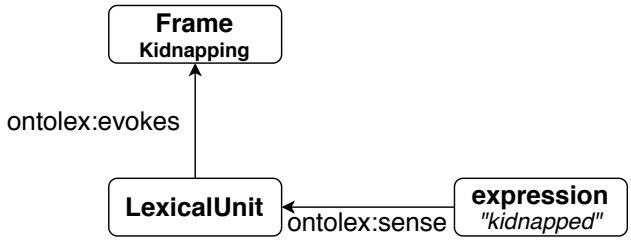

Figure 1: RDF modeling for conceptual relationships. The expression *kidnapped* expresses one of its senses as a lexical unit. This lexical unit evokes the **Kidnapping** frame.

Each frame is further associated with a characteristic set of *frame elements* that apply to the syntactic realization of the phrases dominated by the frame. We refer to Example (1).

(1) **Kidnapping**
[PERPETRATOR Two men] ⊙kidnapped [VICTIM the children] [TIME yesterday].

In this example, 'kidnapped' evokes **Kidnapping**, which consists of several frame elements. 'Two men' expresses the PERPETRATOR frame element and 'the children' expresses the VICTIM frame element. These frame elements are called *core* frame elements, i.e., they need to be overtly specified in order for the reader to process the frame. Other types of frame elements, like 'yesterday', are *peripheral*, meaning that they modify the frame.
When a core frame element is not present in the predicate scope, it is annotated as a *Null Instantiation*, which we paraphrase as being *unexpressed*. In Subsection 3.1., we will elaborate on the phenomenon of unexpressed core frame elements and how we propose to treat them. Frames are situated in semantic space through *frame-to-frame relations*. In these relations, one frame is the more abstract superframe, and the other is the less abstract subframe. One of the relations through which **Kidnapping** is situated has an inheritance frame-to-frame relationship with **Committing_crime**, which is a conceptually corresponding yet less specific superframe . In Subsection 3.1., we will show how frame-to-frame relations are used in FrameNet to explore variation in framing.

2.2. FrameNet annotation tools

To the best of our knowledge, there are four publicly accessible and popular FrameNet annotation tools.
Annotation for English FrameNet (Ruppenhofer et al., 2006) is based on four layers. The *target* is the target word that will be tagged with a frame label. Each constituent of the *target* is a candidate for frame element annotation. Each constituent is labeled with a grammatical function and a phrase type. Grammatical functions are syntactic relations that a constituent fulfills with respect to the target word, e.g., *object* in the case of a verbal target word. Phrase types indicate the syntactic category of the constituent, e.g., *noun phrase* in the case of the constituent 'the man'. English FrameNet annotates one sentence at a time, in which one target word is labeled with a frame and its frame elements. An annotator first labels a target word with a frame label. Consequently, the grammatical function and the phrase type of each of the constituents of the target word are shown, which can be corrected manually. The annotation guidelines are built upon the values of the grammatical function and the phrase type of a constituent, i.e., these notions guide the annotator in deciding which frame element to apply. In the online demo of the annotation tool, the grammatical functions and the phrase types are not shown to avoid visual clutter.
The Global FrameNet Project (Torrent et al., 2018) builds upon the annotation setup of English FrameNet. The core novelty lies in moving to a multilingual setting. The aim is to gain insight into how different languages frame translations of the same texts. This is accomplished by enriching the annotation by allowing annotators to specify why a certain annotation was not possible based on the existing frames, e.g., *too specific* or *too general*.
WebAnno (Eckart de Castilho et al., 2016) is a generic web-based annotation tool for semantic and syntactic structures, of which FrameNet annotation is one of the options. The main emphasis of the tool is on the relation between syntactic and semantic structures, which drives the annotation effort. The tool offers the possibility of introducing con-

[1] We chose to not use the OntoLex (McCrae et al., 2017) relationship *ontolex:reference* since it might lead to confusion in distinguishing between conceptual and referential relationships.

straint rules in order to speed up annotation.
Salto (Burchardt et al., 2006) is a multi-level annotation tool, which can be used to annotate FrameNet information. The annotation starts with a syntactic analysis of a sentence. After determining the target word and labeling it with a frame, the constituents can be tagged with a frame element by means of drag and drop functionality.
All four described annotation tools provide the functionality to annotate FrameNet information. All of them start with a syntactic analysis of the sentence and annotate FrameNet information on top of that analysis. They differ in what syntactic information is used and how this drives the annotations.
The annotation tool described in this paper differs from the existing annotation tools. The main difference stems from the choice of the central element in the tool. Whereas most existing tools use a sentence as the central element, our tool makes use of an incident, i.e., an event, as the primary unit. For each event, the user is shown both structured data and texts that all make reference to the same incident. The purpose of the annotation effort is to make it possible to annotate mentions conceptually by linking to FrameNet, and referentially towards the structured data.

3. Variation in Framing of Reference

In this section, we discuss variation in framing at the referential level. In Subsection 3.1., the means of variation in framing within the FrameNet paradigm is discussed, as well as the merits of adding the referential level. In Subsection 3.2., we introduce a data model to facilitate referential annotations as well as the main data resource used. In Subsection 3.3., we propose to add a relationship between an expression and a frame in order to make the connection between the referent of an expression and its evoked frame explicit.

3.1. Variation of framing in FrameNet

Within FrameNet, variation in framing can be observed by measuring the degree to which different subframes stand in a similar frame-to-frame relation to a superframe. See a classic example below.

(2) a. **Commerce_sell**
[$_{\text{TIME}}$ Yesterday,] [$_{\text{SELLER}}$ John] $_{\odot}$sold [$_{\text{BUYER}}$ Mary] [$_{\text{GOODS}}$ a book].
b. **Commerce_buy**
[$_{\text{BUYER}}$ A woman] $_{\odot}$bought [$_{\text{GOODS}}$ a novel] [$_{\text{PLACE}}$ in the shop].

In (2a), 'sold' evokes **Commerce_sell**, with 'Mary' labeled as the Buyer. In (2b), 'bought' evokes **Commerce_buy**, with 'a woman' labeled as the BUYER. Both frames are related to the abstract frame **Commerce_goods_transfer** and show a different perspective on this event. This way, variation in framing is measured on a conceptual level, comparing different variants of subframes related to one abstract superframe.
In capturing variation in framing at a conceptual level, the annotation provides no knowledge concerning the referential level of the text. For instance, we lack insight as to whether the two predicates in (2) refer to the same event in the real world, which would entail that 'Mary' and 'a woman' refer to the same referent. The current tool aims to implement structured data about the event, enabling the annotator to annotate on both the conceptual and the referential level. This allows us to investigate variation in a broader sense: not just the framing of abstract concepts, but also with respect to the referent.
In addition to variation in subframes belonging to a superframe, variation in framing can be observed when measuring the extent to which core frame elements are expressed. According to FrameNet, core frame elements are necessary components of a frame (Ruppenhofer et al., 2006). Yet, core frame elements often remain unexpressed in a sentence. FrameNet distinguishes between unexpressed core frame elements that are left out due to syntactic constraints or allowances (e.g., passivization, imperatives, pro-drop) and core frame elements that are left out due to anaphoric reasons: they are already given as part of the surrounding context of the sentence. See the examples below, taken from the FrameNet database (Ruppenhofer et al., 2006).

(3) **Change_of_leadership**
a. [$_{\text{NEW\_LEADER}}$ Khan himself] was $_{\odot}$elected [$_{\text{ROLE}}$ a Congress party MP for Rampur].
b. Also [$_{\text{TIME}}$ on July 13] [$_{\text{SELECTOR}}$ the congress] $_{\odot}$elected [$_{\text{NEW\_LEADER}}$ Gorbarev] [$_{\text{FUNCTION}}$ to head a commission [...]]

In both sentences in (3), the verb evokes **Change_of_leadership**. One of the core frame elements of this frame is SELECTOR: the person or group 'responsible for a change in leadership' (FrameNet lexical database; (Ruppenhofer et al., 2006)). In (3a), the SELECTOR is unexpressed, which can be assigned to the syntax, since passivized constructions allow speakers to leave out the agent. However, (3b) shows an active syntax, while the SELECTOR remains unexpressed. Moreover, **Change_of_leadership** contains more core frame elements that are unexpressed in (3), such as, for instance, OLD_LEADER, OLD_ORDER, and BODY. These core frame elements are regarded as part of the contextual knowledge and not considered sufficiently relevant to express.
The current categorization of unexpressed core frame elements in FrameNet is syntax-driven, meaning that these frame elements are analyzed within sentence boundaries. When their absence is assumed to be bounded by sentence-external words or phrases, this information is not further specified. The downside of this approach is that we do not gain insight into the way that these core frame elements are linguistically encoded in the full discourse or if they are encoded at all. Certain approaches address this problem by going beyond the predicate scope in annotating unexpressed core frame elements. For instance, in SemEval-2010 Task 10: Linking Events and Their Participants in Discourse (Ruppenhofer et al., 2010), unexpressed core frame elements were annotated outside of the scope of the predicate in order to gain insight into the referents of these unexpressed roles. A small number of texts from a work of Arthur Conan Doyle were annotated. There were three

participating systems. The results showed that this is a very challenging task for Natural Language Processing systems. Building upon the insights gained from SemEval-2010 Task 10, we consider the text as a cohesive narrative structure and allow for annotation of core frame elements throughout the full text. We hypothesize that frames and frame elements are evoked either directly or indirectly throughout the discourse in relation to the minimally required referential level. The tool, therefore, allows for annotation of frame and frame element relations at both the subword level, e.g., compounds, as well as across sentences. Being able to annotate frames and (core) frame elements throughout the text also allows us to analyze how different sources frame the same situation differently and to explore the underlying factors of unexpression and other differences.
We discuss the implementation of this adaptation in Section 4. and its function in Subsection 5.1.

3.2. Data model & main data resource

To facilitate the combination of conceptual and referential annotation, we make use of a data model in which an incident is the central element.
Let R be a registry of real-world incidents, i.e., event instances. Let R_i be a real-world incident and let $R_i \in R$. Each R_i contains structured data about the real-world incident, e.g., the period or time when the event happened, its location, and information about which participants played a role and in which capacity. Let E_t be an event type, which is a categorization of a real-world incident. Finally, there are reference texts, which are descriptions of real-world incidents (R_i), e.g., a news article describing what happened. Each R_i can have multiple reference texts.
The main data resource used in the annotation tool is Wikidata (Vrandečić and Krötzsch, 2014). We represent a Wikidata item, i.e., a description of an event instance, as an incident. Wikidata provides structured data about the incident, such as the time, location, and participants. Also, a Wikidata item lists Wikipedia pages in multiple languages that make reference to that specific Wikidata item, which we represent as the reference texts. Finally, each Wikidata item is tagged with one or multiple *instance of* (Property P31) relationships, which indicate the event type(s) of the Wikidata item.

3.3. The connection between a frame and a referent of an expression

English FrameNet (Ruppenhofer et al., 2006) uses the *evoke* relationship to relate lexical units to frames. However, the evoke relationship does not provide information about the relationship between the event that an expression refers to and the evoked frame. In our approach, we make this information explicit in our annotation. We clarify the distinction through the examples in Table 1.
Table 1 provides examples that highlight the relationship between the evoked frames and the instances that the expressions refer to. All examples originate from the English Wikipedia page describing the *2006 Hezbollah cross-border raid* (Wikidata identifier Q2026122), in which Hezbollah conducted a raid on Israeli territory in 2006. During the raid, Hezbollah kidnapped Israeli soldiers. In Sentence 1, the target word 'kidnapped' evokes **Kidnapping** and also refers to the event in which Israeli soldiers were kidnapped that is an instance of a *kidnapping* event. Similarly, the noun 'attack' in Sentence 3 evokes **Attack**, and the event it refers to is an instance of the evoked frame. On the contrary, the noun 'kidnappers' in Sentence 2 evokes **Kidnapping**, but it refers to Hezbollah, which means that the instance that the expression refers to is not an instance of the evoked frame but an instance of the concept *person*. These role-designating nouns typically serve as a frame element of the verb they are governed by. Finally, the verb 'can' in Sentence 4 evokes **Possibility**, but it is unclear what it refers to. There is no clear relationship between the evoked frame and what the target word refers to.

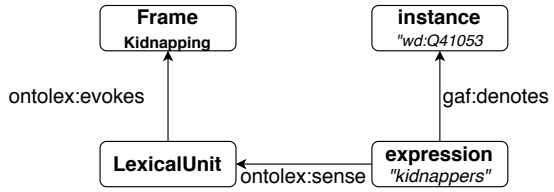

Figure 2: RDF modeling for conceptual and referential relationships without the dfn:isOfFrame relationship. The expression 'kidnappers' refers to the Wikidata item Q41053 (*Hezbollah*) and evokes the **Kidnapping** frame. Since *Hezbollah* is not an instance of kidnapping, no *rdf:type* relationship is assigned.

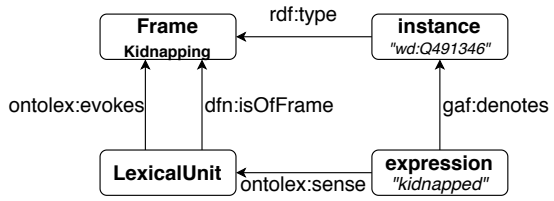

Figure 3: RDF modeling for conceptual and referential relationships with the dfn:isOfFrame relationship. The expression 'kidnapped' refers to the Wikidata item Q491346 (*kidnapping of Kim Dae-jung*) and evokes the **Kidnapping** frame. Since *kidnapping of Kim Dae-jung* is an instance of kidnapping, the *rdf:type* relationship is assigned and hence also the dfn:isOfFrame relationship.

We extend the conceptual RDF relationships with referential ones, for which we use Figures 2 and 3 for clarification purposes. For all target words in Table 1, it is the case that they evoke a frame. For most examples (all except Sentence 4), there is a referential link, which we model via the *gaf:denotes* relationship as part of the GAF framework (Fokkens et al., 2014). In the case that the referent that the expressions refer to is an instance of the evoked frame, we create an *instance of* relationship, for which we use the *rdf:type* relationship, between the incident and the evoked frame (see Figure 3). We make this relationship explicit by establishing a *http://rdf.cltl.nl/dfn/isOfFrame* link between the *LexicalUnit* and the *Frame*. In cases where the referent is not an instance of the evoked frame, the *rdf:type* and *http://rdf.cltl.nl/dfn/isOfFrame* relationship are absent (see Figure 2).

ID	Sentence	POS	Evokes	Refers to	Relation frame to incident
1	Six Western tourists were **kidnapped** by Al-Faran on 4 July 1995.	verb	**Kidnapping**	the kidnapping as part of Wikidata item Q2026122	the referent is an instance of the frame.
2	In December 1995, the **kidnappers** left a note that they were no longer holding the men hostage.	noun	**Kidnapping**	Hezbollah	the referent is not an instance of the frame
3	Top Hezbollah official Ghaleb Awali was assassinated in a car bomb **attack** in the Dahiya in Beirut in July 2004	noun	**Attack**	the car bombing as part of Wikidata item Q2026122	the referent is an instance of the frame.
4	Israel **can** get to Hezbollah anywhere in Lebanon	verb	**Possibility**	-	there is no referential relation

Table 1: Examples sentences taken from the English Wikipedia page describing the *2006 Hezbollah cross-border raid* (Wikidata identifier Q2026122). The first column indicates the example sentence identifier, the second shows the example sentence, the third provides the part of speech tag of the target word, the fourth which frame the target word evokes, the fifth column provides information about what the target word refers to, and the last column indicates the relationship between the evoked frame and its referent.

4. Data-to-text Annotation Tool

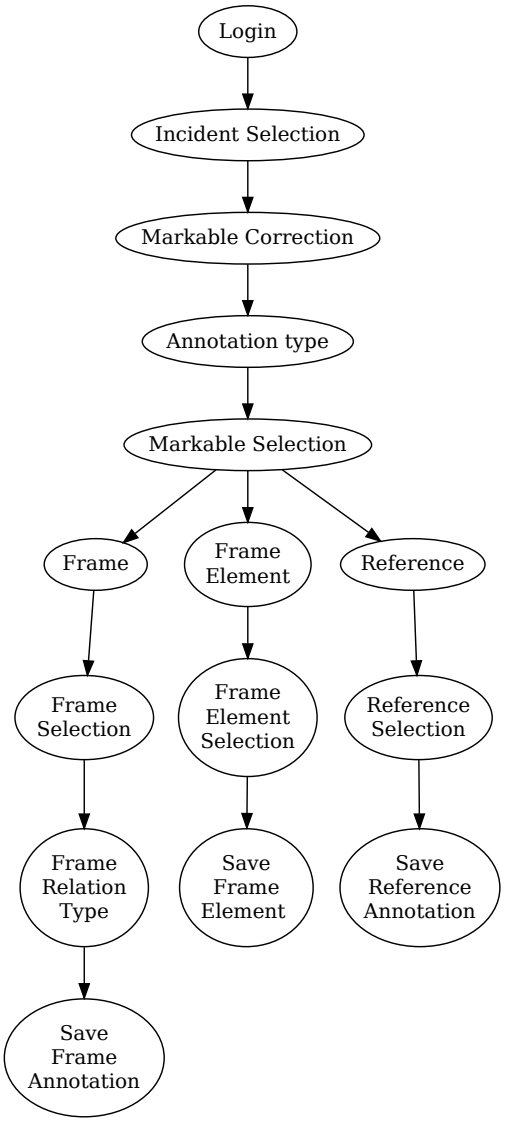

Figure 4: Annotation workflow

In this section, we introduce the annotation tool, for which we present the workflow in Figure 4. The tool starts from data that is aggregated through the MWEP platform described in Vossen et al. (2020). It contains structured data on incidents of a specific type, e.g., *murders*, *elections*, *sports events*, etc. paired with reference texts linked to the specific incidents.

After the login, the annotator first selects an event type, after which a list of incidents is given. Next, the annotator can select a specific incident from the list, after which the structured data is shown with all the reference texts.

The user first has the option to correct the tokenization in the texts to ensure that multi-words and compounds are correctly represented. After deciding on the markables, three types of annotation types can be chosen: *Frame*, *Frame Element*, or *Reference*. With *Frame* and *Frame Element*, the user can annotate predicates with their corresponding frames and frame elements. In contrast, *Reference* is used to link textual mentions to the structured data of the incident. This enables coreferential mentions, i.e., linked to the same incident, to obtain different frame annotations, which forms the basis for analyzing variation.

The front-end of the tool makes use of Bootstrap CSS[2] and jQuery[3], and the server-side operations are handled by Node.js[4].

In Subsection 4.1., we introduce the resources used in the tool. The Subsections 4.2., 4.3., 4.4., 4.5., 4.6., 4.7., and 4.8. explain the main components of the annotation tool.

4.1. Resources

In this subsection, we introduce the resources used in the annotation tool, i.e., the lexicon and the data.

lexicon We make use of the canonical version 1.7 of FrameNet (Fillmore and Baker, 2010; Ruppenhofer et al., 2006). All annotations make use of a Resource Description Framework (RDF) of FrameNet, for which the two most common resources are Framester (Gangemi et al., 2016)

[2]https://getbootstrap.com/docs/3.4/css/

[3]https://jquery.com/

[4]https://nodejs.org/en/

and PreMOn (Corcoglioniti et al., 2016). We chose to use PreMOn since the project was more active.[5]

data acquisition We have developed a data architecture (Vossen et al., 2020) to obtain and represent the data according to the data model as presented in Subsection 3.2., for which we primarily make use of Wikidata (Vrandečić and Krötzsch, 2014).

preprocessing spaCy[6] is used for sentence splitting, tokenization, and part of speech tagging, for which models in English, Dutch, and Italian are used. The preprocessing is stored in the NLP Annotation Format (NAF) (Fokkens et al., 2014), a stand-off, multilayered annotation schema for representing linguistic annotations.[7] We retrained OpenSESAME (Swayamdipta et al., 2017) to tag the reference texts with FrameNet frames.[8]

4.2. Login

A unique session identifier is created for each annotator for each annotation session. Each annotation will then be accompanied by this session identifier and the timestamp of the annotation, which allows analyses per annotation session and per annotator. No annotations are removed. Automatically generated annotations are represented in the same way using identifiers and timestamps.

4.3. Incident Selection

The user will first have to choose a specific data release, e.g., *version 1.0*. From this data release, an event type is chosen, e.g., *murder* (Wikidata identifier *Q132821*). From the available incidents that belong to the chosen event type, one incident is selected.

After clicking on *Load Incident*, the user is presented with the structured data about the incident, e.g., the location, time, and participants. Also, all available reference texts that make reference to the selected incident are shown, possibly in multiple languages.

The existing annotations for each reference text are highlighted. The user can observe the difference between manual and automatic annotations, which is designed such that the user can focus more on validating than on full-text annotation.

4.4. Markable Correction

Linguistic phenomena in which there is a many-to-many relationship between a token and a concept are a crucial problem for language technology (Sag et al., 2002). Idioms, phrasal verbs, and compounds are cases in which this occurs. In at least two phenomena, i.e., idioms and phrasal verbs, multiple tokens combined refer to one concept or semantic unit (Lexicon of Linguistics, 2020b; Quirk, 2010). In contrast, compounds consist of one token, but they can evoke multiple frames and frame element relations.

In this step of the annotation process, the user can correct the automatic tokenization by indicating which combination of tokens serve as phrasal verbs or idioms. Also, the user can decompose compounds into separate components.

For cases in which multiple tokens should be merged, the user clicks on the tokens that are part of the construction and indicates whether they belong to the category of phrasal verbs or idioms. We follow English FrameNet in assigning the part of speech tag *V* to phrasal verbs and *idio* for idioms. If the user now clicks on one of the tokens of a construction, all tokens that belong to it are selected. It is no longer possible to annotate parts of the construction as predicates or frame elements. Also, any annotation on the level of the individual parts of the construction is deprecated and will no longer be used nor rendered in the tool.

Annotators are also asked to detect *endocentric* compounds, i.e., compounds consisting of a grammatical head and a modifier (Lexicon of Linguistics, 2020a), which make it possible to annotate components of the compounds with frames and frame elements. After clicking on the detected endocentric compound, the user is asked to indicate the components of the compound as well as which component serves as the frame-evoking unit. For each component, the user needs to indicate the lemma and the part of speech according to the Universal Dependencies version 2 (Nivre et al., 2017) part of speech tagset. After specifying a compound, the user can now click on the separate components of the compound and can no longer click on the compound as a whole. Also, any previous annotation of the compound as a whole is ignored and will no longer be rendered.

4.5. Annotation Type

The next step involves deciding which type of annotation to perform. There are three options: *Frame* (see Subsection 4.6.), *Frame Element* (see Subsection 4.7.), and *Reference* (see Subsection 4.8.). For each annotation, this is the first step. Note that our tool does not assume that there is already a FrameNet lexicon beforehand.

4.6. Frame

The goal of the *Frame* annotation type is to annotate predicates with FrameNet frames as well as to indicate the *Frame Relation Type*. After selecting a markable, the user clicks to observe a dropdown list in which all FrameNet frames are divided into four groups: 1. **Typical frames** this category contains the frames that are typically expected given the type of the selected incident, e.g., **Killing** and **Offenses** for the event type *murder*. 2. **candidate frames for lemma and part of speech**: the candidate frames given the lemma and part of speech of the markable are shown here. 3. **candidate frames for lemma**: the candidate frames given the lemma of the markable are shown here. 4. **other** all other FrameNet frames are shown here.

When a user selects a frame from the dropdown list, more information about the frame is shown in the right panel.

Also, the user has to indicate the *Frame Relation type* (see Subsection 3.3.). In the case that the incident to which the expression refers is an instance of the evoked frame, the user selects *isOfFrame*. Otherwise, the user selects *evoke*.

[5] We downloaded the dataset from the following link: `https://knowledgestore.fbk.eu/files/premon/dataset/latest/premon-2018a-fn17-noinf.tql.gz`.

[6] https://spacy.io/

[7] `https://github.com/newsreader/NAF`

[8] Our wrapper is available at: `https://github.com/cltl/run_open-sesame`.

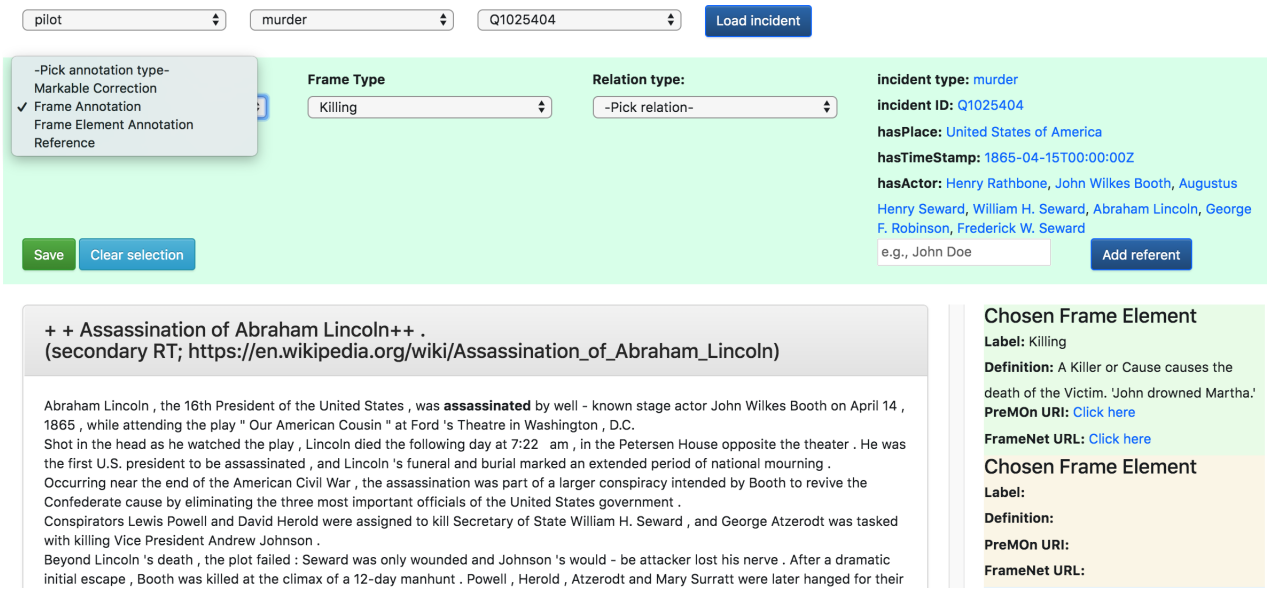

Figure 5: Snapshot of the data-to-text annotation tool.

After clicking *Save*, the annotation is added to the corresponding NAF file of the reference text using the PreMOn URI identifier for the frame.

Finally, we allow the user to click multiple predicates at once and annotate a batch of predicates with the same frame label and frame relation type.

4.7. Frame Element

The goal of the *Frame Element* Annotation Type is first to indicate which frame elements are found in the predicate scope. If core frame elements are not found in the scope of the predicate, the user is asked to try to annotate them in the context surrounding the predicate.

Given that a user has previously annotated a predicate with a frame label, the user can now also annotate frame elements for this predicate. The user clicks on the markable and selects the frame element from a dropdown list. After annotating at least one frame element for a frame, e.g., KILLER for **Killing**, a table is shown in the right panel of the tool, of which an example is shown in Table 2.

Frame Element	Type	Annotated	Expressed
KILLER	Core	true	true
VICTIM	Core	false	false
CAUSE	Core	false	false
MEANS	Core	false	false
INSTRUMENT	Core	false	false

Table 2: Frame Element Annotation

Table 2 presents the information shown to the user during the frame element annotation phase. The user can keep track of all frame element annotations for an active frame. The annotator is asked to attempt to find evidence for each core frame element in the predicate scope. If the frame element is unexpressed, the user is asked to look for evidence of the frame element in the surrounding context, e.g., the VICTIM is mentioned in the sentence before the target sentence. Only after all core frame elements have been annotated, the user is able to switch to a different Annotation Type.

After clicking *Save*, the annotation is added to the corresponding NAF file of the reference text using the PreMOn URI identifier for the frame element.

4.8. Reference

Alongside the reference texts, the user is shown structured data about the main incident, as shown in the top right corner of Figure 5. The structured data table consists of five pieces of information. The *event type* of the incident is shown, which we obtain from the Wikidata *instance of* relationship (Property P31) as well as the Wikidata item identifier. Also, we categorize the properties of a Wikidata item into three classes and model them using the Simple-Event-Model (SEM (Van Hage et al., 2011)). Locations are mapped to sem:hasPlace, temporal expressions to sem:hasTimeStamp, and participants to sem:hasActor.

The user clicks on some tokens and then indicates, by clicking on a link in the structured table, that the markable refers to the reference, e.g., 'The kidnapper' refers to Hezbollah (Wikidata item Q41053) which is expressed in RDF using a gaf:denotes relation.

The annotator can also modify the structured data table. He or she can add and remove values to sem:hasPlace, sem:hasTimeStamp, and sem:hasActor, provided that the user provides Wikidata items as values.

5. Discussion

In this section, we elaborate on the ways in which the current tool supports us in providing frame annotations of texts referring to a real-world event. In Subsection 5.1., we discuss the extent to which the tool in its current state directs the annotator. In Subsection 5.2., we present future plans for the tool to capture inferred frames.

ID	Sentence	Evokes	Frame Element
1	*The Tunisian man* who prosecutors say **perpetrated** last month's terrorist attack [...]	**Committing_crime**	PERPETRATOR
2	*he* **ploughed** a truck into a crowded Christmas market	**Impact**	AGENT
3	*Amri* **carried out** the attack	**Intentionally_act**	AGENT
4	*he* **hijacked** a truck	**Piracy**	PERPETRATOR
5	*he* [...] **shot** its Polish driver	**Hit_target**	AGENT
6	*he* [...] **drove** it into the crowded market	**Operate_vehicle**	DRIVER

Table 3: Example sentences taken from reference texts referencing *Anis Amri* (Wikidata identifier Q28052669), the agent of the *2016 Berlin Attack* (Wikidata identifier Q28036573). The first column indicates the identifier, the second column shows the example sentence with the reference to Anis Amri in italics and the frame-evoking predicate in bold, the value of the third column is the evoked frame, and the last column shows which frame element the reference to Anis Amri expresses.

5.1. Functionality

The current annotation tool enables the annotator to perform two parallel annotations. In choosing Frame or Frame Element as the Annotation Type, the annotator traditionally performs frame annotations. In addition, he or she can select the Reference Annotation Type to mark words that refer to the structured data. These parallel annotations can be performed on the same expression in the text, which means that this expression is annotated as both contributing to a specific frame and simultaneously referencing a component of the real-world event. The resulting annotation scheme displays which words refer to the structured data as well as how these words frame the data. In other words, variation in framing across texts can be measured concerning a fixed real-world referent. For instance, all sentences in Table 3 make reference to the 2016 Berlin attack. In these sentences, *Anis Amri* is being referred to by the expressions 'The Tunisian man', 'he', and 'Amri'. However, these expressions belong to different predicates, each evoking a different frame. Hence, each expression is labeled with a different frame element, which provides insight into how the referent is framed.

A second function of the tool is to annotate unexpressed core frame elements beyond the scope of the predicate. Given the assumption that frames construct a cohesive narrative in referring to a main event, we expect these core frame elements to occur somewhere in the reference text to support the reader's understanding of the narrative. If no evidence is found for a core frame element, this raises questions about its retrievability: whether it is inferred from world knowledge or expressed in other reference texts. Example 4 highlights this functionality.

(4) **Change_of_leadership**

a. [ROLE Presidential] ⊙elections were held [LOCATION in Slovakia] [TIME in March 2019].

b. [OLD_LEADER Incumbent President Andrej Kiska] did not run for a second term.

The sentences in (4) form the onset of a text referring to the presidential election of Slovakia in 2019. Assuming that the user first annotates Example (4a), 'elections' is marked as a predicate evoking **Change_of_leadership**, which contains many core frame elements, which are ROLE, BODY, FUNCTION, NEW_LEADER, OLD_LEADER, OLD_ORDER, and SELECTOR. Example (4a) contains annotations for three frame elements, which are Role, Location, and Time. Role is the only core frame element with an annotation in Example (4a) out of the many core frame elements. No evidence is found for the other core frame elements. The next step for the user is to find mentions of these unexpressed core frame elements. Evidence for the specific old leader of the election, i.e., core frame element OLD_LEADER, is found in Example (4b). The annotator continues to look for evidence for the remaining core frame elements.

Finally, we provide the annotator with a list of typical frames. The list consists of frames that are relevant for the perception of the event, e.g., **Offenses** and **Use_firearm** for a *murder* event. Some of those typical frames are even necessarily evoked, e.g., **Killing** in the case of a *murder* event. The main rational to use this list is to restrict the annotations to the most important mentions of the event, based on the assumption that only a subset of frames is used to describe an event of a specific type. Moreover, if the obligatory typical frames are not found in a text, this then leads to questions about evocation through inference.

5.2. Desired functions

During preliminary annotation experiments with the tool, we found that a substantial portion of the frames indicated by the typical frames are not derived from reference texts. For instance, **Killing** is not evoked by lexical units in a reference text of a murder event. We argue that the necessary frames are still activated in the text, but that they are derived through pragmatic inference rather than lexical evocation. These inferences are derived from linguistic cues that are not marked as lexical units within FrameNet. Hence, these are different inferences than the ones discussed by Chang et al. (2002), who use the notion of inference in FrameNet for a frame that is inevitably processed during the evocation of another frame, e.g. **Commerce_buy** is always activated with the evocation of **Commerce_sell** and vice versa. These inferences are actually based on a frame that is evoked by a lexical unit. Also, most inferences we detected could not

be derived by any current frame-to-frame relations. See the sentence in (5), which refers to the *2016 Berlin Attack* (Wikidata identifier Q28036573).

(5) **Killing**
[$_{\text{CAUSE}}$ a truck] was $_{\odot?}$deliberately $_{\odot?}$driven into the Christmas market, $_{\odot?}$leaving [$_{\text{VICTIM}}$ 12 people] $_{\odot?}$dead

Although the words in (5) separately do not evoke **Killing**, the sum of the components 'driven', 'leaving', and 'dead' activate this frame by means of entailment. Moreover, 'deliberately' acts as a cue from which **Offenses** could be derived. The way these frames are activated can only be traced by complementing the lexical semantic analysis of frame semantics with a pragmatic analysis in terms of inference. If the annotator is guided in pointing out the linguistic cues, the different ways in which frames are inferred from these cues can be schematized. The data can then be analyzed with respect to their pragmatic type, e.g., entailment, implicature (Levinson, 1983); and the possible factors that account for the inference e.g., historical distance of the publication or genre conventions.

One of the next steps in the development of the annotation tool is to implement an *inferred frame layer* for the annotator that allows him or her to mark linguistic cues from which a frame is pragmatically derived, on top of the traditional FrameNet annotation module. After targeting words as lexical units, this inferred frame layer will ask the annotator to mark *n* linguistic cues in the text that might derive any of the remaining unannotated typical frames. We refer to Remijnse and Minnema (2020) for a detailed description of this proposal.

6. Conclusion

In this paper, we introduced an annotation tool in which annotations can be made to both the conceptual and referential level. For an event type, the tool delivers a collection of incidents, each accompanied with structured data and reference texts in different languages. The annotator can mark targets in the texts to frame-annotate and mark the same targets to annotate referential relations to the structured data. From the output of this annotation process, patterns of variation in framing can be extracted concerning reference to a single referent.

7. Acknowledgements

The research reported in this article was funded by the Dutch National Science organisation (NWO) through the project *Framing situations in the Dutch language.*

8. Bibliographical References

Andor, J. (2010). Discussing Frame Semantics: The State of the Art: An Interview with Charles J. Fillmore. *Review of Cognitive Linguistics. Published under the auspices of the Spanish Cognitive Linguistics Association*, 8(1):157–176.

Burchardt, A., Erk, K., Frank, A., Kowalski, A., Pado, S., and Pinkal, M. (2006). SALTO-A Versatile Multi-Level Annotation Tool. In *Proceedings of the 5th International Conference on Language Resources and Evaluation (LREC 2006)*.

Chang, N., Petruck, M. R. L., and Narayanan, S. (2002). From Frames to Inference. In *Proceedings of the First International Workshop on Scalable Natural Language Understanding*.

Corcoglioniti, F., Rospocher, M., Aprosio, A. P., and Tonelli, S. (2016). PreMOn: a Lemon Extension for Exposing Predicate Models as Linked Data. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 877–884, Portorož, Slovenia, May. European Language Resources Association (ELRA).

Cybulska, A. and Vossen, P. (2011). Historical Event Extraction from Text. In *Proceedings of the 5th ACL-HLT Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities*.

Eckart de Castilho, R., Mújdricza-Maydt, É., Yimam, S. M., Hartmann, S., Gurevych, I., Frank, A., and Biemann, C. (2016). A Web-based Tool for the Integrated Annotation of Semantic and Syntactic Structures. In *Proceedings of the Workshop on Language Technology Resources and Tools for Digital Humanities (LT4DH)*.

Fillmore, C. J. and Baker, C. (2010). A Frames Approach to Semantic Analysis. In *The Oxford handbook of linguistic analysis*. Oxford University Press.

Fokkens, A., Soroa, A., Beloki, Z., Ockeloen, N., Rigau, G., Van Hage, W. R., and Vossen, P. (2014). NAF and GAF: Linking Linguistic Annotations. In *Proceedings 10th Joint ISO-ACL SIGSEM Workshop on Interoperable Semantic Annotation*.

Gangemi, A., Alam, M., Asprino, L., Presutti, V., and Recupero, D. R. (2016). Framester: A Wide Coverage Linguistic Linked Data Hub. In *European Knowledge Acquisition Workshop*.

Levinson, S. C. (1983). *Pragmatics*. Cambridge Textbooks in Linguistics. Cambridge University Press.

Lexicon of Linguistics. (2020a). Exocentric compound. `https://lexicon.hum.uu.nl/?lemma=Endocentric+compound&lemmacode=778&lemma=Endocentric+compound&lemmacode=778`.

Lexicon of Linguistics. (2020b). Idiom. `https://lexicon.hum.uu.nl/?lemma=Endocentric+compound&lemmacode=778&lemma=Endocentric+compound&lemmacode=778`.

McCrae, J. P., Bosque-Gil, J., Gracia, J., Buitelaar, P., and Cimiano, P. (2017). The Ontolex-Lemon model: development and applications. In *Proceedings of eLex 2017 conference*.

Nivre, J., Agić, Ž., Ahrenberg, L., Aranzabe, M. J., Asahara, M., Atutxa, A., Ballesteros, M., Bauer, J., Bengoetxea, K., Bhat, R. A., Bick, E., Bosco, C., Bouma, G., Bowman, S., Candito, M., Cebiroğlu Eryiğit, G., Celano, G. G. A., Chalub, F., Choi, J., Çöltekin, Ç., Connor, M., Davidson, E., de Marneffe, M.-C., de Paiva, V., Diaz de Ilarraza, A., Dobrovoljc, K., Dozat, T., Droganova, K., Dwivedi, P., Eli, M., Erjavec, T., Farkas, R., Foster, J., Freitas, C., Gajdošová, K., Galbraith, D., Garcia, M., Ginter, F., Goenaga, I., Gojenola, K., Gökırmak, M., Goldberg, Y., Gómez Guinovart, X., Gonzáles Saave-

dra, B., Grioni, M., Grūzītis, N., Guillaume, B., Habash, N., Hajič, J., Hà Mỹ, L., Haug, D., Hladká, B., Hohle, P., Ion, R., Irimia, E., Johannsen, A., Jørgensen, F., Kaşıkara, H., Kanayama, H., Kanerva, J., Kotsyba, N., Krek, S., Laippala, V., Lê H`ông, P., Lenci, A., Ljubešić, N., Lyashevskaya, O., Lynn, T., Makazhanov, A., Manning, C., Mărănduc, C., Mareček, D., Martínez Alonso, H., Martins, A., Mašek, J., Matsumoto, Y., McDonald, R., Missilä, A., Mititelu, V., Miyao, Y., Montemagni, S., More, A., Mori, S., Moskalevskyi, B., Muischnek, K., Mustafina, N., Müürisep, K., Nguy˜ên Thị, L., Nguy˜ên Thị Minh, H., Nikolaev, V., Nurmi, H., Ojala, S., Osenova, P., Øvrelid, L., Pascual, E., Passarotti, M., Perez, C.-A., Perrier, G., Petrov, S., Piitulainen, J., Plank, B., Popel, M., Pretkalniņa, L., Prokopidis, P., Puolakainen, T., Pyysalo, S., Rademaker, A., Ramasamy, L., Real, L., Rituma, L., Rosa, R., Saleh, S., Sanguinetti, M., Saulīte, B., Schuster, S., Seddah, D., Seeker, W., Seraji, M., Shakurova, L., Shen, M., Sichinava, D., Silveira, N., Simi, M., Simionescu, R., Simkó, K., Šimková, M., Simov, K., Smith, A., Suhr, A., Sulubacak, U., Szántó, Z., Taji, D., Tanaka, T., Tsarfaty, R., Tyers, F., Uematsu, S., Uria, L., van Noord, G., Varga, V., Vincze, V., Washington, J. N., Žabokrtský, Z., Zeldes, A., Zeman, D., and Zhu, H. (2017). Universal Dependencies 2.0. LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.

Quirk, R. (2010). *A Comprehensive Grammar of the English Language*. Pearson Education India.

Remijnse, L. and Minnema, G. (2020). Towards Reference-Aware FrameNet Annotation. In *The International FrameNet Workshop 2020*.

Ruppenhofer, J., Ellsworth, M., Schwarzer-Petruck, M., Johnson, C. R., and Scheffczyk, J. (2006). *FrameNet II: Extended theory and practice*. FrameNet Project.

Ruppenhofer, J., Sporleder, C., Morante, R., Baker, C., and Palmer, M. (2010). SemEval-2010 Task 10: Linking Events and Their Participants in Discourse. In *Proceedings of the 5th International Workshop on Semantic Evaluation*.

Sag, I. A., Baldwin, T., Bond, F., Copestake, A., and Flickinger, D. (2002). Multiword Expressions: A Pain in the Neck for NLP. In *Computational Linguistics and Intelligent Text Processing*.

Swayamdipta, S., Thomson, S., Dyer, C., and Smith, N. A. (2017). Frame-Semantic Parsing with Softmax-Margin Segmental RNNs and a Syntactic Scaffold. *arXiv preprint arXiv:1706.09528*.

Torrent, T. T., Ellsworth, M., Baker, C., and Matos, E. (2018). The Multilingual FrameNet Shared Annotation Task: a preliminary report. In *Proceedings of the 11th International Conference on Language Resources and Evaluation (LREC 2018)*.

Van Hage, W. R., Malaisé, V., Segers, R., Hollink, L., and Schreiber, G. (2011). Design and use of the Simple Event Model (SEM). *Web Semantics: Science, Services and Agents on the World Wide Web*, 9(2):128–136.

Vossen, P., Ilievski, F., Postma, M., and Roxane, S. (2018). Don't Annotate, but Validate: a Data-to-Text Method for Capturing Event Data. In *11th edition of the Language Resources and Evaluation Conference (LREC 2018)*.

Vossen, P., Ilievski, F., Postma, M., Fokkens, A., G., M., and L., R. (2020). Large-scale Cross-lingual Language Resources for Referencing and Framing. In *12th Edition of the Language Resources and Evaluation Conference (LREC 2020)*.

Vrandečić, D. and Krötzsch, M. (2014). Wikidata: a Free Collaborative Knowledge Base. *Communications of the ACM*, 57(10):78–85.

9. Language Resource References

Burchardt, A., Erk, K., Frank, A., Kowalski, A., Padó, S., and Pinkal, M. (2009). FrameNet for the Semantic Analysis of German: Annotation, Representation and Automation. *Multilingual FrameNets in Computational Lexicography: methods and applications*, 200:209–244.

Djemaa, M., Candito, M., Muller, P., and Vieu, L. (2016). Corpus Annotation Within the French FrameNet: a Domain-by-domain Methodology. In *The 10th International Conference on Language Resources and Evaluation (LREC 2016)*.

Gruzitis, N., Nespore-Berzkalne, G., and Saulite, B. (2018). Creation of Latvian FrameNet based on Universal Dependencies. In *Proceedings of the International FrameNet Workshop*.

Hayoun, A. and Elhadad, M. (2016). The Hebrew FrameNet Project. In *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC 2016)*.

Heppin, K. F. and Gronostaj, M. T. (2012). The Rocky Road towards a Swedish FrameNet-Creating SweFN. In *The 8th international conference on Language Resources and Evaluation (LREC 2012)*.

Laviola, A., Lage, L., Marção, N., Tavares, T., Almeida, V., Matos, E., and Torrent, T. (2017). The Brazilian Portuguese Constructicon: Modeling Constructional Inheritance, Frame Evocation and Constraints in FrameNet Brasil. In *2017 AAAI Spring Symposium Series*.

Ohara, K. H., Fujii, S., Ohori, T., Suzuki, R., Saito, H., and Ishizaki, S. (2004). The Japanese Framenet Project: An Introduction. In *Proceedings of LREC-04 Satellite Workshop "Building Lexical Resources from Semantically Annotated Corpora"(LREC 2004)*.

Ruppenhofer, J., Ellsworth, M., Schwarzer-Petruck, M., Johnson, C. R., and Scheffczyk, J. (2006). *FrameNet II: Extended theory and practice*. FrameNet Project.

Subirats, C. and Sato, H. (2003). Surprise! Spanish FrameNet. In *Proceedings of the Workshop on Frame Semantics at the XVII. International Congress of Linguists*.