Legal Knowledge and Information Systems
E. Francesconi et al. (Eds.)

doi:10.3233/FAIA220470

WhenTheFact: Extracting Events from European Legal Decisions

María NAVAS-LORO [1] and Víctor RODRÍGUEZ-DONCEL
Ontology Engineering Group, Universidad Politécnica de Madrid, Madrid, Spain

Abstract. This paper presents WhenTheFact, a tool that identifies relevant events from European judgments. It is able to extract the structure of the document, as well as when the event happened and who carried it out. WhenTheFact builds then a timeline that allows the user to navigate through the annotations in the document.

Keywords. event extraction, visualization, NLP, legal domain, timeline generation

1. Introduction

Events and their logical sequence are key to understanding legal decisions, being the storyline of pivotal importance. We therefore assume that a judgment can be described as a series of time-marked happenings (*events*) instead of focusing on the other entities (things), and to this aim we must be able to extract these events in an automatic fashion.

Before undertaking the event extraction task itself, discourse extraction is required; since legal decisions are long and complex, where the event is detected within the document is crucial regarding its relevance. Once the relevant parts of the document are determined, the next step involves training a system using documents annotated manually with relevant events, as well as the semantic resources available. Finally, the system is able to annotate different documents, allowing to visualize the relevant events in it. Additionally, in the online demo[2], a timeline with these relevant events is generated, easing navigation through the document.

The paper is organized as follows. Section 2 explores previous related work in literature. Section 3 introduces the system created to extract relevant events from legal decisions, explaining its different stages: document structure extraction, training strategies and extraction itself. Section 4 presents the evaluation of the system. Finally, Section 5 summarizes the main contributions and the future research lines to explore.

2. Related work

Beside generic efforts in event extraction such as the carried out by temporal taggers following TimeML [1,2] or related tasks such as frame-semantic parsing [3,4], semantic

---

[1]Corresponding Author: María Navas-Loro, Ontology Engineering Group, Universidad Politécnica de Madrid, Spain; E-mail: mnavas@fi.upm.es. This work was funded partially by the project Knowledge Spaces: Técnicas y herramientas para la gestión de grafos de conocimientos para dar soporte a espacios de datos (Grant PID2020-118274RB-I00, funded by MCIN/AEI/10.13039/501100011033) and by H2020 MSCA PROTECT (813497).

[2]https://whenthefact.oeg.fi.upm.es/

role labeling (SRL) [5,6] or open information extraction[3], some proposals have been made specifically in the legal domain. These works often involve *ad hoc* definitions of events, ignoring general event annotation schemes.

In the context of legal information retrieval, events can be considered as temporally bounded objects that have entities important as participants that played a significant role in a case. To this aim, Lagos et al. [7] propose an NLP semi automatic approach to enable the use of entity related information corresponding to the relations among the key players of a case, extracted in the form of events. They are interested in the topic, the roles, the location and the time, and consider different types of events. On the other hand, Maxwell et al. [8] reviewed 150 events extracted from 18 sentences from the Canadian Supreme court and compared them with automatic extraction using SRL on two cases. Another approach was done for Spanish [9], looking for patterns in documents that help them identify legal events and related information (*who, what, to whom* and *where*), and analyzing the verbs that occur in the texts. In order to improve information retrieval in Brazilian courts, similar work was performed for Portuguese [10].

In summary, legislation systems consist still of semiautomatic or even manual approaches. Most of the proposals within the legal domain tend to be supported by patterns, using manually crafted rules or semantic role labelling techniques [8,7].

3. Event Extraction

Based on a previous works about temporal expressions in the legal domain [11], first step for building a knowledge graph was to decide the source of the documents, since there are important differences among jurisdictions, even when they share the language. Due to the ease of importing and reusing judgments from their respective repositories, as well as the multilingual challenge it offers and the possible associated documents that could eventually be added to a knowledge graph, we decided to work with decisions from European courts, namely the European Court of Human Rights (ECHR) and European Court of Justice (ECJ). Choosing a specific source also allowed us to analyze the structure of the documents, which improves the ability to extract relevant events [12]. Regarding the format of the annotations, we will use the one specified in the EventsMatter corpus [13][4].

The remaining of this section will present the structure extractor of the judgments (Section 3.1, the different training strategies used (Section 3.2) and the pipeline of the event extraction algorithm (Section 3.3), that applies the two previous techniques.

3.1. Structure Extraction

From an analysis performed in the EventsMatter [13], the only available corpus of judgments annotated with events (to the best of our knowledge), we can confirm the importance of the sections in identifying which events are relevant and which are not. To this end, we have developed a Structure Extractor that

1. Detects the structure of the document and divides it into parts with a title, a type, a parent and the begin and end offsets.

[3] https://stanfordnlp.github.io/CoreNLP/openie.html
[4] A very preliminary version of this work was briefly introduced in the paper describing this corpus.

2. Looks for the most relevant sections and sends the sentences within to the algorithm that extracts the events, ignoring sections such as references to laws.

This Structure Extractor is currently able to handle the structure of the ECHR and ECJ documents, but in such a way that a new document type can be easily added. Additionally, if for any reason the processed documents did not adhere to the expected structure (for example, with very old cases that followed a different format), it would simply return all sections.

3.2. *Training Strategies*

Regarding the training strategy of the event extraction system, we used both semantic and syntactic considerations. On the one hand, we collected all the events and attached arguments annotated in the training set of the EventsMatter corpus [13]. The EventsMatter corpus is a collection of 30 legal decisions manually annotated with events and their arguments (namely, *who*, *when* and *what*, called *core*). Once collected, we stored both the core of the events and the relations among their different parts. On the other hand, we also used an external semantic resource, FrameNet, to enrich the keywords we use to identify legal events. Subsequent sections provide a detailed description of both approaches.

3.2.1. *EventsMatter Training Set*

The first step of the training phase was to collect all the event mentions in the corpus training set. We isolated then the parts of the sentences annotated as core and generated a sentence just with it, adding has generic subject "They" in order to make them simple to parse and grammatically correct. Thereupon we iterated over all these *simple* sentences, creating a frame for each of the main verbs of the sentences that stored the information of all the mentions of this verbs along the corpus. This is, that for instance the verb "lodge" (that is to some extent a *light verb*[5] in the legal domain) can appear in several sentences carrying different meaning depending on the object attached. Some examples of its use would be the constructions "lodge a complaint", "lodge a request", "lodge an appeal", "lodge an objection" or "lodge an action". It should be noted that most of these cases could be simplified using a single semantic-carrying verb, such as "to complain" or "to request", but that the legal domain tends to recur to these paraphrasing in texts, since they usually imply not just an action but also a formal procedure (usually administrative).

Each of verbs found in this phase constitutes a *frame* that will be used to identify and classify future mentions of each of the verbs in new texts. Finally, it must be noted that we distinguish between passive and active voice when searching for the dependency parsing relations among the members of the core of an event. This is a consideration that might not be important in general kind of texts, but the legal domain tends to present a high rate of passive verbs. Among the events in the training set, for instance, we find that the 14% of the mentions were expressed as passive sentences.

Two couples of text files containing (1) the *simple* version of each sentence with a relevant event mention and (2) the type of events of each of the mention are available

[5] *Light verbs* are those verbs that have little semantic meaning, needing therefore more words to constitute a full predicate. This is for instance the case of the verbs "make" or "take" in English. For more information on this linguistic phenomenon, please check the work by Butt [14].

within the system – a couple for all the sentences of the corpus (named *all*) and another for just the training part (*train*). The collection of events can be easily extended by adding to the files new sentences and their respective types, and a detailed example of this Frame structure can be found in the website.

3.2.2. FrameNet training

It is straightforward that some events not present in the training set of the EventsMatter corpus should be detected in other documents, and even that events considered not relevant in those documents can be relevant in other cases.

This is why, in addition to the events gathered from the training set explained previously, we decided to enrich the system with frames from FrameNet [15]. FrameNet is a database that contains semantic frames together with the words that represent them in text, as well as additional information such as the arguments this frame can present. Since frames represent situations, they can be understood as events to some extent, and incorporating a selection of them to our target events would help to generalize our approach. Since not all the frames in FrameNet are of interest, we manually inspected the database using the FrameGrapher tool[6], that allowed us to navigate through it and find the most relevant frames to our task. After examining the different relations among the frames, we found the most general ones, as well as their children, and imported their information. These most legally representative parent frames were namely "Committing_crime", "Crime_scenario", "Law", "Obligation_scenario", and "Misdeed". The frames collected from them, together with the lexical units associated to them (that is what we will look for in the text), are detailed in the webpage. A text file containing this information is available in the system. In order to add more frames, it is only necessary to add them to the file maintaining the same format.

3.3. Event Extraction

Regarding the event extraction itself, Fig. 1 depicts the pipeline of the tool. We detail the different stages of the processing below.

First step consists of finding the relevant parts of the text to annotate, using for this the Structure Extractor detailed in Section 3.1. If the structure is not recognized, the whole text will be annotated, what obviously impacts in a negative way in the amount and quality of the events. Otherwise, just the relevant parts of the document are processed subsequently.

Next step is to find the sentences involving temporal expressions. To this aim we adapt and integrate the functionality of Añotador [16], a temporal tagger able to recognize temporal expressions. If there is at least one temporal expression in a sentence, we check if it is a special case (namely the application lodgement, that always follows the same syntactic structure). If so, we annotate the arguments and go to the next sentence. If not, we check if the sentence contains any of the events gathered from the training corpus. If so, we do the dependency parsing (*deppar*) of the sentence (using CoreNLP [17]) and check if it is valid and look for the arguments (see (1) below). If not, we check again for the legal frames specifically selected from FrameNet. If this is the case, we check them similarly that in the events case (see (2)). Once we detected the main event

[6] https://framenet.icsi.berkeley.edu/fndrupal/FrameGrapher

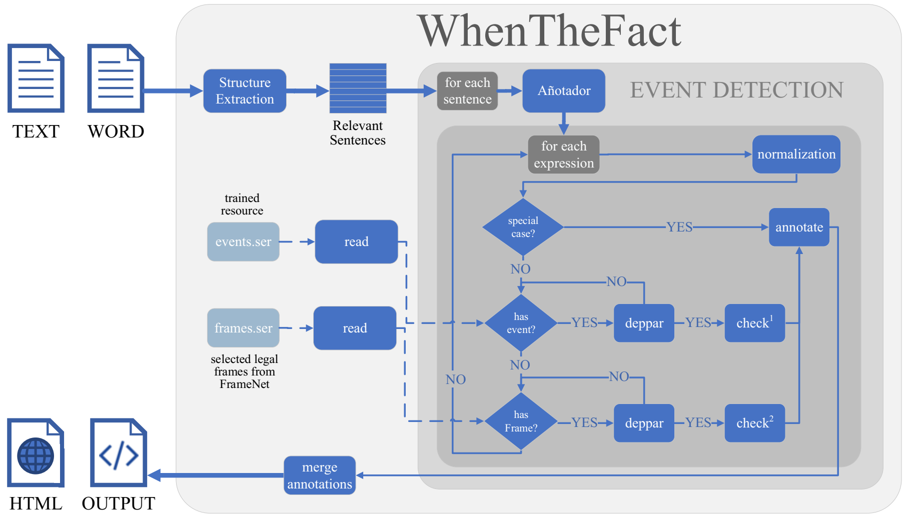

Figure 1. Pipeline of WhenTheFact.

in the sentence, if there was more than one temporal expression in it, we will select the temporal expression that is the closest to the core of the event.

(1) For the events, we check if it is not an auxiliary verb nor in the gerund form. Then we check if it is in passive or active voice. Depending on this, we will look either for the relations gathered from passive training cases or from active ones.
(2) For the frames, the check function is similar to the events' one, but there are no specific relations stored for each frame, so the argument "who" and the extent of the core are therefore detected using default relations.

Once all the sentences have been explored, we merge all the annotations and produce the output. This output consists of an annotated XML and as a visual HTML that also includes a timeline built from the retrieved events.

4. Evaluation

Regarding evaluation, we have compared WhenTheFact's results against the EventsMatter corpus and checked it has improved. The evaluation is depicted in Table 1.

5. Conclusions

In this paper we have presented WhenTheFact, an event extractor able to annotate relevant legal events taking into account the structure of a legal judgment. Next steps include solving correference, currently uncovered, since for now we just get the textual mention, that can consist of pronouns. Once this is achieved, queries will be able to retrieve for instance the timeline of one actor's involvement in a case.

Table 1. Comparison between the previous implementation of the WhenTheFact event extractor (OLD) and the new implementation (NEW).

		Event				**Event Components**					
		Identification		Type		What		When		Who	
		Len	Str	Len	Str	Len	Str	Len	Str	Len	Str
OLD	P	85.71	80.00	47.14	42.86	80.26	23.68	77.59	72.41	75.00	68.75
	R	77.92	72.73	42.86	38.96	69.32	20.45	63.38	59.15	63.16	57.89
	F	81.63	76.19	44.90	40.82	74.39	21.95	69.77	65.12	68.57	62.86
NEW	P	86.49	81.08	54.05	51.35	83.75	82.50	79.03	74.19	81.43	74.29
	R	83.12	77.92	51.95	49.35	76.14	29.55	69.01	64.79	75.00	68.42
	F	84.77	79.47.19	52.98	50.33	79.76	30.95	73.68	69.17	78.08	71.23

Also multilinguality is currently being explored. Although several approaches have been tested already, none of them has been good enough to guarantee acceptable results for all the languages. Finally, the tool can be used not just for visualization, but also to populate legal knowledge graphs to be used in different contexts.

References

[1] Verhagen M, et al. Automating Temporal Annotation with TARSQI. In: Proceedings of the ACL 2005 on Interactive Poster and Demonstration Sessions. ACL; 2005. p. 81-4.
[2] Chambers N, et al. Dense event ordering with a multi-pass architecture. Transactions of ACL. 2014;2:273-84.
[3] Swayamdipta S, et al. Frame-Semantic Parsing with Softmax-Margin Segmental RNNs and a Syntactic Scaffold. arXiv preprint arXiv:170609528. 2017.
[4] Alam M, et al. Semantic role labeling for knowledge graph extraction from text. Progress in Artificial Intelligence. 2021:1-12.
[5] Gardner M, et al. AllenNLP: A Deep Semantic Natural Language Processing Platform; 2017. .
[6] Agerri R, et al. IXA pipeline: Efficient and Ready to Use Multilingual NLP tools. In: LREC. vol. 2014; 2014. p. 3823-8.
[7] Lagos N, et al. Event extraction for legal case building and reasoning. In: International Conference on Intelligent Information Processing. Springer; 2010. p. 92-101.
[8] Maxwell KT, et al. Evaluation of semantic events for legal case retrieval. In: Proceedings of the WSDM'09 Workshop. ACM; 2009. p. 39-41.
[9] Sierra G, et al. Event Extraction from Legal Documents in Spanish. In: Workshop on Language Resources and Technologies for the Legal KG, LREC; 2018. .
[10] Bertoldi A, et al. Cognitive Linguistic Representation of Legal Events. Towards a semantic-based legal information retrieval. In: COGNITIVE 2014; 2014. .
[11] Navas-Loro M, et al. TempCourt: evaluation of temporal taggers on a new corpus of court decisions. The Knowledge Engineering Review. 2019;34:e24.
[12] Navas-Loro M, Santos C. Events in the legal domain: first impressions. In: Proceedings of the 2nd Workshop TeReCom (JURIX 2018); 2018. p. 45-57.
[13] Filtz E, et al. Events Matter: Extraction of Events from Court Decisions. In: Proceedings of JURIX. 336; 2020. p. 33-42.
[14] Butt M. In: The light verb jungle: still hacking away. Cambridge University Press; 2010. p. 48–78.
[15] Baker CF, Fillmore CJ, Lowe JB. The Berkeley FrameNet Project. In: Proceedings of ACL/COLING - Volume 1. ACL; 1998. p. 86–90.
[16] Navas-Loro M, Rodríguez-Doncel V. Annotador: a temporal tagger for Spanish. Journal of Intelligent & Fuzzy Systems. 2020;39:1979-91. 2.
[17] Manning CD, et al. The Stanford CoreNLP Natural Language Processing Toolkit. In: Proceedings of the 52nd Annual Meeting of the ACL, System Demonstrations; 2014. p. 55-60.