



Towards Grammatical Tagging for the Legal Language of Cybersecurity

Gianpietro Castiglione
gianpietro.castiglione@phd.unict.it
University of Catania, Department of
Mathematics and Computer Science
Italy

Giampaolo Bella
giamp@dmi.unict.it
University of Catania, Department of
Mathematics and Computer Science
Italy

Daniele Francesco Santamaria
daniele.santamaria@unict.it
University of Catania, Department of
Mathematics and Computer Science
Italy

ABSTRACT

Legal language can be understood as the language typically used by those engaged in the legal profession and, as such, it may come both in spoken or written form. Recent legislation on cybersecurity obviously uses legal language in writing, thus inheriting all its interpretative complications due to the typical abundance of cases and sub-cases as well as to the general richness in detail. This paper faces the challenge of the essential interpretation of the legal language of cybersecurity, namely of the extraction of the essential Parts of Speech (POS) from the legal documents concerning cybersecurity.

The challenge is overcome by our methodology for POS tagging of legal language. It leverages state-of-the-art open-source tools for Natural Language Processing (NLP) as well as manual analysis to validate the outcomes of the tools. As a result, the methodology is automated and, arguably, general for any legal language following minor tailoring of the preprocessing step. It is demonstrated over the most relevant EU legislation on cybersecurity, namely on the NIS 2 directive, producing the first, albeit essential, structured interpretation of such a relevant document. Moreover, our findings indicate that tools such as SpaCy and ClausIE reach their limits over the legal language of the NIS 2.

KEYWORDS

Privacy, Data Protection, Act, Pronouncement, NLP, POS tagging

ACM Reference Format:

Gianpietro Castiglione, Giampaolo Bella, and Daniele Francesco Santamaria. 2023. Towards Grammatical Tagging for the Legal Language of Cybersecurity. In *The 18th International Conference on Availability, Reliability and Security (ARES 2023)*, August 29–September 01, 2023, Benevento, Italy. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3600160.3605069>

1 INTRODUCTION

It is widely accepted that, in its attempt to fully describe all applicable circumstances and details, legal language may become lengthy and complicated, hence various professional figures specialise precisely in the interpretation of the contents of documents written in

such a language. Legal language is, of course, also used at the level of European legislation, which promotes several types of legal acts, namely Regulations, Directives, Decisions, Recommendations and Opinions [2].

European legislation also touches upon cybersecurity, notably through the “Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union”, known as NIS 2 in brief [6], which can be considered the most relevant treatment of cybersecurity in legal language at European level today. We observe that lengthy written information may complicate, in general, even the essential interpretation step known as *grammatical tagging* or *POS tagging*, which consists of extracting the main components of the discourse, such as clauses of a sentence and subjects of a clause. This might be problematic not only for professionals such as lawyers but also for the layman who is called to read and accept various terms of service on the web several times a day.

In light of these observations, we want to investigate how to use computer-support to assist through the POS tagging of legal language. The present work leverages NLP techniques to develop a methodology offering machine-assisted support through the extraction of the Parts Of Speech (POS) of legal language. In consequence, the understanding of directives and similar documents does no longer have to be entirely manual but can be automated by applying our methodology. More precisely, the task our methodology automates is information extraction for POS tagging, which consists in the identification of the main agents and of the actions they must perform according to the given text.

Therefore, our methodology seeks out to extract at least the relevant clauses from each sentence as well as to identify the relevant POS from each clause, namely the subject (agent), the main verb or predicate (action), and the object of the action (object). Because we aim at open science, our work leverages the forefront of open-source techniques, namely the SpaCy and ClausIE libraries, as we shall see below. We understood from our early experiments that, once tailored to legal language, such libraries reach their limits, hence our methodology requires the analyst to validate their outputs by a systematic comparison with a manual analysis of sentences.

Our methodology is arguably general, hence it can be applied to any legal document. However, we specifically demonstrate it on the legal language of modern cybersecurity directives, in particular on the NIS 2 Directive, which can be considered the best benchmark. We analysed its 46 articles and selected those where actual prescriptions are made, leaving away those about general prerequisites and terminology. As a result, we grammatically analysed articles 7 through to 37 by both a parser based on ClausIE and manually,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ARES 2023, August 29–September 01, 2023, Benevento, Italy

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0772-8/23/08...\$15.00

<https://doi.org/10.1145/3600160.3605069>

concluding that manual validation cannot be avoided over legal language. Precisely, we found that the tools are good in tagging subjects on average on 96.1% of cases, then verbs on 83.3% and objects only on 64.5% of cases.

An overview of the related work and of SpaCy are in Section 2 and in Section 3 respectively. The method for grammatical tagging of legal language is illustrated step by step in Section 4 while its application to NIS 2 Directive is in Section 5. The statistical analysis is evaluated in Section 6 and the open challenges encountered are highlighted in Section 7. In Section 8 the conclusions are outlined.

2 RELATED WORK

Several works have been conducted in the legal field by leveraging NLP and machine learning algorithms. The new capabilities made available by artificial intelligence have even allowed the prediction of judgements, the search for inconsistencies between the various legal processes and other tools for the simplification of legal context. This Section only outlines the most recent and most closely related to the contributions of this paper.

The LEGAL-BERT tool adapts the BERT models to the legal context; the main feature is to predict some deliberately maskable words related to the legal context [5]. The scope of the prediction is even wider; Medvedeva et al. investigated the use of NLP for predicting legal decisions, and in particular for predicting if there were any violations of the European Convention on Human Rights [14]. At the same time, Katz et al. constructed a model, based on the random forest approach, for predicting, using only available data, years and years of decisions by the Supreme Court of the United States [10].

Particular attention has been given to the use of NLP in the context of patents: Arts et al. use “natural language processing techniques to harness the rich content of patent documents, identify new technologies and their impact... validation studies support the use of text mining techniques to identify new technologies and measure patent novelty at the time of filing, and to measure the impact of these new technologies on subsequent innovation.” [3]; while Sheng Lee et al. developed a deep learning pre-trained model for the generation of patent claims [11].

NLP techniques have been not used only for prediction purposes and for text generation but also for questioning and answering. Zhong et al. presented a question-answering data-set for the legal domain that collects questions from the National Judicial Examination of China [20]. They reason on the data-set for facing the challenges of word matching, concept understanding or multi-paragraph reading. Similar classifications for legal documents were built by de Araujo et al. with the additional feature of leveraging NLP for theme assignment and labelling, regarding Brazilian Supreme Court [13], and by Sulea et al., which apply machine learning techniques for predicting and investigating on the sentences of French Supreme Court [17].

Artificial intelligence techniques have been used also to solve particular cases of inconsistencies. Xu et al. developed a framework for legal judgment prediction tasks [19]. The goal of the framework is to face confusing charges by analysing similar articles and fact descriptions, thanks to a new graph distillation operator. Ul Hassan et al. leverage NLP and machine learning algorithms for legal text

classification [18]. They introduced new models for identifying and extracting, from construction contracts, the two categories of requirements and non-requirements. Lippi et al. face the cases of terms of service of online platforms. In particular, they leveraged machine learning for detecting if any unfair clauses for the user of the online platform exist [12].

In the mere context of information extraction two main branches can be identified: POS tagging and topic modelling. In the first case, several libraries for different programming languages exist, for example, Scikit-learn [16]; in the second case, Latent Dirichlet Allocation (LDA) is one of the most used approaches [9]. LDA, introduced by Blei, Ng and Jordan [4], is a generative probabilistic model and represents documents as random mixtures over several topics, where a topic is produced by distribution among different words.

It seems fair to claim, from the summary given above, that the main works are based on ad-hoc development of NLP or machine learning algorithms within a specific use case. In particular, the main techniques aim at the prediction of text and situations or, alternatively, at classification as their main goals. Such approaches are certainly valid but bear the inherent limitations of being heavily dependent on pre-trained models related to the specific use case, which in turn require the availability of huge amounts of labelled data and consequently high-performance hardware.

3 AN OUTLINE OF SPACY

Following the related work just summarised, it is clear that POS tagging is only a specific niche of NLP and that, in particular, no significant attempts to apply POS tagging to legal language are worth noting. Following these observations, we aim at approaches that are lightweight and, at the same time, general so as to be applicable to virtually any legal language. In the context of the work presented in this paper, we shall use what is perhaps the best-established tool for processing text, namely the SpaCy library [8], whose main peculiarity beyond POS is Named Entity Recognition.

SpaCy is a free library for the Python programming language and can be considered among the most versatile tools to tailor NLP to practical applications. In order to get the correct token, SpaCy assigns syntactic dependency labels that allow us to identify subjects, verbs or objects. SpaCy both works on pre-trained pipelines related to a specific language hence context-independent, as well as on local data. The SpaCy pipeline of actions is essentially as follows: first, the words in input are tokenised. After that, each token is tagged to a specific part of speech, leveraging the pre-trained models, thereby predicting the closest role of each token derived from the input.

In practice, it is very useful to appeal to a sub-library of SpaCy called ClausIE (claucy) [7], which is able to divide the input text into the main components of each statement, namely into clauses. We shall see below several examples supporting the case that the complexity of a legal text can be daunting even for such powerful, state-of-the-art NLP tools. Moreover, we are aware that POS tagging is only one step towards automated information extraction from such a language, while its semantic interpretation can be expected to be even more challenging.

4 A METHOD FOR GRAMMATICAL TAGGING OF LEGAL LANGUAGE

Our method for the grammatical tagging of any document written in legal language follows a waterfall style and is completed with manual validation, as we shall see below. In short, given a legal document, it takes as input some text written in a legal language and carries out a punctual analysis of the input in order to output the POS tagging of each part of the text.

Step 1. Preprocessing. In order for our method to also automate document parsing, a preliminary preprocessing phase is necessary. Our target, in fact, is not just information extraction but includes automating the selection of each sentence subjected to tagging. Since each legal document has its own structure, the parsing must necessarily be manually preset and context-dependent. In general, we want to accomplish the following steps:

a) Identification of articles. Assuming we are analysing the legal language of a document that is composed of articles, this step consists in identifying the more significant ones where we can extract relevant POS. These are normally those articles where the actions to be carried out by the entities defined in the document itself are defined in precise detail. By contrast, our target articles do not include those with abbreviations and general provisos. It follows that the correct articles can only be determined with a manual and subjective selection of the range to be considered. In fact, if we parse all the articles without distinction, we could get an inconsistent collection of the POS, because some might not represent the notions of actions we aim to distil.

b) Text acquisition. After finding the range of articles to examine, this step consists of acquiring all the text from a given Article with the aim of automated processing. We shall expect a subdivision into items, which can be numeric or alphabetical, within each article, a style that is quite common in legal documents. A useful acquisition heuristic is to select all text between the string “Article *i*” and “Article *i+1*” in order to collect all text of the *i*-th article, of course having checked that such keywords do not occur in the text forming the body of the article. Moreover, if the selection involves two different pages, then spurious information must be deleted.

c) Item selection. To input a text, such as a legal one, to any NLP technique, it is necessary to define a heuristic for the choice of the *single units of processing* inside the given text, namely the fragments of text that, given as a whole to the chosen technique, leads to producing new, relevant information. By contrast, applying it to a whole article, which may consist of several sentences, may not lead to useful information. As our aim is to identify single subjects, actions and objects, our single unit of processing is represented by each single sentence.

d) Hierarchy identification. Sometimes, the single unit of processing may have to be built by copying parts of the statement, for example, subject and action, which may have been omitted to limit redundancy. This is typically the case when an item from an itemised list is specialised by a whole itemised sub-list. In such a typical case, this step associates the underlying item with the overlying one, namely by removing the colon and linking the underlying hierarchical level with the one above. However, it is necessary to

manually verify the unit of processing that is built this way so as to minimise the risk of inconsistencies.

Step 2. POS tagging. As mentioned, grammatical tagging of a sentence extracts the parts of speech and, first of all, identifies the clauses from each unit of processing. Once clauses are identified, each clause may contain up to 5 different sentence patterns: Subject (S), Verb (V), Complement (C), Object (O), Adverb (A). In particular, the difference between C and O relies on the role of each of them: a Complement completes the meaning of a sentence while Object is specifically the direct object on which the action of the verb is reflected. Of course, while only one S and V pair should be expected per clause, several C and O instances may be encountered.

For example, patterns can be extracted by the code snippet Code 1, which leverages our chosen libraries. It can be seen that, firstly, it is necessary to load the pre-trained model, and the choice of *en_core_news_lg* means adopting the largest model consisting of 514k keys for a total of 560 MB. In all steps of the method, this represents the highest computational cost but it is arguably necessary to deal with as many cases as possible. The pre-trained model is then loaded into the pipeline and the “engine” for text pattern extraction is ready. The code returns the extracted clauses together with all sentence patterns.

Code 1: Python code snippet for the extraction of sentence clauses and patterns based on SpaCy and ClausIE

```
import spacy
import claucy

nlp = spacy.load("en_core_news_lg")
claucy.add_to_pipe(nlp)

def NLP(text):
    document = nlp(text)
    return document._.clauses,
        document._.clauses[0].to_propositions(as_text=True)
```

The same thing can be programmed in a different way, as illustrated in Code 2 [15]. The main difference is that Code 2 leverages SpaCy only. For the object extraction, in fact, the dependency tree generated by the tagging step will be parsed. Only when the tag related to the object is obtained, the sub-text identified by the related sub-tree will be returned. This can be limiting and less precise than using ClausIE, since the latter uses a search based on sentences rather than tags. Our preliminary experiments support this claim, in fact, it could happen that the objects identified by ClausIE would not be identified by SpaCy alone, indicating that our first approach is preferable to the second, while performance difference is negligible.

Code 2: Python code snippet for the extraction of sentence clauses and patterns based on SpaCy only

```
def get_object_phrase(doc):
    for token in doc:
        if ("dobj" in token.dep_):
            stree = list(token.subtree)
            a = stree[0].i
```

```
b = stree[-1].i + 1
return doc[a:b]
```

Step 3. Tabulation. For a correct statistical analysis of the validity of the method, relevant data must be collected and structured. The template that will be filled in for each single article is shown in Table 2. For each relevant element we extracted by NLP functions in the previous step we dedicate three columns, for a total of 10 columns considering an additional one for the identification of the single item. In particular, the first column features the correct POS that we extracted by a manual analysis of the sentence, then the second column states the POS extracted by our use of the NLP libraries, while the third column indicates an evaluation of the correctness of the automated extraction with respect to our manual one, using an alphabet of four symbols, represented in Table 1. In particular, a tick symbol means that the two analyses coincide, then a cross and a vertical bar symbol respectively mean that we found the automated analysis to be entirely wrong or partially wrong. Additionally, a “P” symbol stands for clauses in passive style.

Table 1: Legend of symbols

Symbol	Acronym
✓	Correct answer
✗	Wrong answer
!	Correct but incomplete answer
P	Passive Verb

The last row of Table 2 refers to the percentage of validity of the solutions found by the automated approach. It could be calculated by a simple metric, for example by assigning full value to the ticks, no value to the crosses, half value to the bars and no value to passive clauses.

5 DEMONSTRATING OUR METHOD ON THE NIS 2 DIRECTIVE

This section illustrates the application of our method to the NIS 2 Directive. For the sake of demonstration, the method is illustrated on Article 11 and Article 23 because they cover different cases. For each step, the following extracts will be analysed:

2. Member States shall ensure that their CSIRTs jointly have the technical capabilities necessary to carry out the tasks referred to in paragraph 3. Member States shall ensure that sufficient resources are allocated to their CSIRTs to ensure adequate staffing levels for the purpose of enabling the CSIRTs to develop their technical capabilities.
 ...
 8. At the request of the CSIRT or the competent authority, the single point of contact shall forward notifications received pursuant to paragraph 1 to the single points of contact of other affected Member States.

Step 1. Preprocessing. By coherently applying the various sub-steps, we have the following outcomes:

a) Identification of articles. In the NIS 2 Directive, the articles are numerated from 1 to 46. However, not all of them are significant from the point of view of the measures. In fact, all articles until 7 are preliminary considerations and definitions, while from the 38 onward various considerations on the applicability of the Directive

are described. Therefore, we will parse the Directive from Article 7 to Article 37.

b) Text acquisition. To acquire the text of Article *i*, the more efficient solution in the context of NIS 2 Directive is as anticipated above, namely to grab the text between each pair of article headers of the form “Article *i*” and “Article *i+1*” as strings. Moreover, we must make sure to expand each header with two characters of line feed both at the start and at the end because some articles could be cited through some items both by the numerical notation and by their full title, a scenario that would hinder the parsing. By contrast, stating which and where the line feed characters are ensures that we capture an occurrence at the beginning of an article, namely an actual header.

An emblematic case in which the described heuristic is necessary is represented by the following extract:

Article 10
Computer security incident response teams (CSIRTs)

1. Each Member State shall designate or establish one or more CSIRTs. The CSIRTs may be designated or established within a competent authority. The CSIRTs shall comply with the requirements set out in Article 11(1), shall cover at least the sectors, subsectors and types of entity referred to in Annexes I and II, and shall be responsible for incident handling in accordance with a well-defined process.
 ...
 10. Member States may request the assistance of ENISA in developing their CSIRTs.

Article 11
Requirements, technical capabilities and tasks of CSIRTs

c) Item selection. Coherently with the definition of this step, a single sentence is selected by referring to its termination by means of the full stop, in the simplest case. The errors that could arise through this approach concern the case in which the full stop does not terminate the sentence. This case only happens in the case of hierarchical lists and is rather frequent through the NIS 2. It will be treated below.

d) Hierarchy identification. In the NIS 2 Directive, each item could have up to two hierarchical levels of depth. The first is listed by letters, while the second is listed by ordinal numbers. Developing a parser that identifies such cases is not complex, since bulleted lists can be used as selection criteria. By contrast, it may be difficult to understand what is represented at the lowest hierarchical level. In the simplest case, we can resolve the following form of compression (Article 9, item 4):

That plan shall lay down, in particular: (a) the objectives of national preparedness measures and activities;

by making the preliminary part of the statement explicit in all possible items, for example, as follows:

That plan shall lay down, in particular, the objectives of national preparedness measures and activities;

In such a case, it generally happens that the primary sentence expresses the presence of some tasks that are described at the first hierarchical level. Consequently, the first hierarchical level states significant verbs that describe the specific actions of the defined

Table 2: Table template for grammatical tagging of an article

N.	Sub	I-Sub	Sub-HIT	Verb	I-Verb	Verb-HIT	Obj	I-Obj	Obj-HIT
...
	Sub-HIT % = %			Sub-HIT % = %			Obj-HIT % = %		

Table 3: Table for grammatical tagging of NIS 2 Article 11

N.	Sub	I-Sub	Sub-HIT	Verb	I-Verb	Verb-HIT	Obj	I-Obj	Obj-HIT
1	-	-	-	-	-	-	-	-	-
1.a	C	C	✓	ensure	ensure	✓	N11.1.2.a	S11.1.2.a	I
1.b	C Premises and SYS	C premises	I	P - located	located	✓	N11.1.2.b	S11.1.2.b	X
1.c	C	C	✓	P - equipped	equipped	✓	N11.1.2.c	S11.1.2.c	X
1.d	C	C	✓	ensure	ensure	✓	N11.1.2.d	S11.1.2.d	✓
1.e	C	C	✓	P - staffed	ensure ensure	✓	N11.1.2.e	S11.1.2.e	X
1.f	C	C	✓	P - equipped	equipped	✓	N11.1.2.f	S11.1.2.f	X
2.1	MS	MS	✓	ensure	ensure	✓	N11.2.1	S11.2.1	✓
2.2	MS	MS	✓	ensure	ensure	✓	N11.2.2	S11.2.2	✓
...
	Sub-HIT % = 95 %			Sub-HIT % = 90 %			Obj-HIT % = 45 %		

subject. Therefore, the second hierarchical level aims at further completing the meaning, defining the objects that the action affects.

Step 2. POS tagging. By executing Code 1 on some sentences of the previous extracts, the following outputs are produced:

Code 3: NLP output on Article 11

```
[<SVC, Member States, shall ensure, None, None, that their CSIRTS jointly have the technical capabilities necessary to carry out the tasks referred to in paragraph 3, []>, <SVO, their CSIRTS, have, None, the technical capabilities necessary to carry out the tasks referred to in paragraph 3, None, [jointly]>]
```

Code 4: NLP output on Article 23

```
[<SV, the single point of contact, shall forward, None, None, None, []>, <SV, the single point of contact shall forward notifications, received, None, None, None, [At the request of the CSIRT or the competent authority, pursuant to paragraph 1, to the single points of contact of other affected Member States]>]
```

In the present work, the desired combinations of sentence patterns are SVO or SVC because in these combinations there are all the POS that are necessary for a subsequent semantic interpretation. If complex sentences are analysed, the combinations given as output may not be unique and definitive, as in the cases above, where two patterns were extracted for each article. Therefore, it would be necessary to manually choose the most appropriate pattern.

In consequence, obtaining only pattern SV may be taken to signify incompleteness of the automated approach, while pattern SVOC falls into the cases that can be considered complete. The whole output consists of sentence pattern type, sentence subject,

sentence verb, sentence indirect object, sentence direct object, sentence complement and sentence adverbials. Our experiments indicated that all possible combinations of the 5 can be output on the NIS 2.

It can be seen from Code 3 that two solutions were generated, both containing the object, but only the first solution is complete. In fact, the correct Subject was identified, together with the Verb and the Complement. Although the second identified sentence pattern might have seemed reasonable, the POS identified are not correct.

Code 4 shows that the produced sentence pattern SV indicates no object was identified. While we can see that both identified sentence patterns produced the correct Subject, the second sentence pattern is wrong in the handling of the verb. This further confirms that the automated approach is not perfect.

A significant case that was encountered is the presence of different hierarchical levels with a specific feature: the object of the sentence introducing the hierarchical level explicitly expresses their presence. This leads to inconsistencies. For example: *Member State shall have the following tasks: a) ... b) ...*. There is greater complexity towards grammatical tagging in the presence of the second hierarchical level.

Step 3. Tabulation. Iterating the extraction over all relevant Articles allows us to collect all information we aim at. The POS of articles 11 and 23 are depicted respectively in Table 3 and Table 4, while the used acronyms are defined in Table 5. Our symbols purposely differentiate the incomplete answer from the incorrect answer, since the first one is to be considered partially correct. Therefore its extraction does not have to be considered entirely misleading.

The tables exhaustively list the occurrences of relevant extracted POS. The tables provide a clear and first-level view of the NLP results previously discussed. Naturally, more correct results are obtained where there is less complexity of the related POS, therefore in most of the cases, simplicity concerns the subjects.

Table 4: Table for grammatical tagging of NIS 2 Article 23

N.	Sub	I-Sub	Sub-HIT	Verb	I-Verb	Verb-HIT	Obj	I-Obj	Obj-HIT	
1.1	MS	MS	✓	ensure	ensure	✓	N23.1.1	S23.1.1	✓	
1.2	Entities concerned	Entities Concerned	✓	ensure	ensure	✓	N23.1.2	S23.1.2	✓	
1.3	MS	MS	✓	ensure	ensure	✓	N23.1.3	S23.1.3	✓	
1.4	Mere ... notification	Mere ... notification	✓	not subject	subject	✗	N23.1.4	S23.1.4	✓	
...	
8	POC	POC	✓	forward	NONE	✗	N23.8.1	S23.8.1	✗	
9.1	POC	POC	✓	submit	submit	✓	N23.9.1	S23.9.1	✓	
9.2	E	E	✓	contribute	contribute	✓	N23.9.2	S23.9.2	✓	
9.3	E	E	✓	inform	inform	✓	N23.9.3	S23.9.3	✓	
10	C	C	✓	provide	provide	✓	N23.10.1	S23.10.1	✗	
11	Commission	Commission	✓	adopt	adopt	✓	N23.11.1	S23.11.1	✗	
			Sub-HIT % = 92.8 %				Sub-HIT % = 80.9 %			
							Obj-HIT % = 83.3 %			

Table 5: Legend of acronyms

Name	Acronym
Member State	MS
National Cybersecurity Strategy	NCS
Competent Authority	CA
Commission	Co
Point of Contact	POC
CSIRT	C
ENISA	E
Cooperation Group	CG
The European External Action Service	EEAS
CN	CSIRTs Network
Eu-Cyclone	EuC
SYS	Supporting Information Systems

6 STATISTICAL ANALYSIS

To evaluate the overall functioning of the automated tagging with respect to our manual one, we average the hit values for each tag of each grammatical tagging table. As suggested above, we associate hit value 1 with a correct answer, 0.5 with an incomplete answer and 0 otherwise. Therefore, for each table i and for each tag j , the Hit Rate (HR) is the sum of the HIT values divided by the number of occurrences:

$$HR_{i,j} = \frac{\sum HV_{i,j}}{\#occ_{i,j}} \cdot 100, \text{ where } 7 \leq i \leq 37, 1 \leq j \leq 3.$$

We decided to skip the following cases: measures with no item number, measures that may include a priori contextualisation and hierarchical levels introduced by “have following tasks” and similar ones where no other agent is expressed. The reasons are, respectively: difficulty in automatically parsing the measure since it has no identifier (for example, Article 17 and Article 36); information extraction is sometimes wrong because the libraries focus on the contextualisation clause; the extracted information is often wrong because from the moment the clause is compressed with the higher hierarchical level, inconsistencies are formed (except in the case where the additional level contains its own subject that allows generalising to the main case of subject, verb and object). Therefore, in cases that have structural rather than semantic problems, since the results would always be wrong and would have consistently

lowered the percentage, we preferred to consider them exceptions. The resulting percentages are shown in Table 6.

As can be seen, the highest values are obtained over Subject and Verb. It may be argued that Objects tend to have lower HIT rates because they may be more complicated. The cases in which they are rather good, compared to the verb of the same item, on the other hand, are due to the presence of hierarchical levels that involve a greater number of sub-cases whose summation raises the value.

7 OPEN CHALLENGES

As demonstrated by the statistical analysis, challenges can be encountered, with a certain relevance, through all cases of information extraction.

Identification of the Object. The challenge of this type does not stem from the complexity of a single word: the pre-trained model is very good at identifying the English words encountered in the text. This can be demonstrated with a single and simple part of speech tagging.

On the contrary, it is difficult for the NLP library to extract the clauses and associate them with the correct place in the sentence. An improvement here would be a breakthrough since one could remove the sentences that simply lengthen the sentence, without adding much additional information, and keep only those sentences that provide significant information.

Associating vast portions of text to either Complement (C) or Object (O) as it currently works limits extraction twice: even irrelevant details are associated causing an overly general tag and, at the same time, the specificity of the individual clauses is lost because many clauses are often tagged together.

Identification of the style. An article may begin with a lengthy contextualisation preamble. This is not infrequent in legal language and serves to contextualise the specific conditions for the application of the very measure that is described in the sequel of the article. This turns out to be a really challenging case for the automated approach. In fact, the POS are often extracted from the early parts in the sentence, with the result that they are obtained from the preamble rather than from the actual measure in this case, hence they are incorrect.

Table 6: Hits of the automated approach wrt the manual one

Art.	Sub-HIT	Verb-HIT	Obj-HIT
7	100 %	100 %	91.6 %
8	100 %	88.8 %	61.1 %
9	100 %	100 %	43.3 %
10	100 %	93.3 %	60 %
11	95 %	90 %	45 %
12	100 %	57.1 %	58.3 %
13	85.7 %	85.7 %	50 %
14	90 %	100 %	50 %
15	95 %	80 %	43.8 %
16	100 %	87.5 %	66.6 %
17	-	-	-
18	100 %	50 %	75 %
19	100 %	83.3 %	57.5 %
20	75 %	75 %	100 %
21	90 %	60 %	60 %
22	100 %	100 %	50 %
23	92.8 %	80.9 %	83.3 %
24	80 %	80 %	40 %
25	100 %	100 %	50 %
26	100 %	33.3 %	33.3 %
27	100 %	75 %	83.3 %
28	100 %	91.6 %	86.6 %
29	100 %	100 %	66.7 %
30	100 %	100 %	100 %
31	100 %	85.7 %	42.9 %
32	100 %	75 %	67.1 %
33	93.8 %	75 %	90 %
34	90 %	70 %	65 %
35	100 %	100 %	50 %
36	-	-	-
37	100 %	100 %	100 %
Average	96.1 %	83.3 %	64.5 %

Grammar issues. Among these, we find the management of passive verbs. In their presence, the object is almost never identified. In fact, managing passive contexts is still a strong limitation for SpaCy. Another problem is the use of pronouns. Of course, this is not exclusively connected to NLP but the association of nouns with their respective pronouns certainly remains an open problem, especially when referring to subjects.

8 CONCLUSIONS

This paper contributed to the essential interpretation of documents written in legal language by advancing an automated methodology for their grammatical tagging. The methodology leveraged the state of the art in the area of NLP, yet prescribing that the various outputs be validated by manual analysis. The methodology can be considered a general finding of this paper, namely it is arguably applicable to any target legal document following some preprocessing tailored to the specific features of the target.

The application of our methodology to the NIS 2 leads to the second finding, which consists of the first, albeit essential, structured

interpretation of the entire directive. We have built all tables for the grammatical tagging of articles from 7 through to 37, showing the POS produced by leveraging the NLP libraries next to those produced manually by our own analysis and, finally, the relevant hit rates [1]. Altogether, these tables may be leveraged in various ways in the future, for example as a reference even by those exercising the legal profession, and by ontologists aiming at the semantic representation of the directive.

A more technical finding is that forefront libraries SpaCy and ClauseIE are very powerful and offer remarkable help. However, they are clearly put at stake by the complexities of the legal style. As a result, while their hit rates are very good over subjects and good over verbs, they are inadequate over objects. Therefore, our future work includes a more in-depth tailoring of these libraries, perhaps by re-training their models specifically for legal documents. We may expect the field of the automated processing of legal language to achieve valuable results in the near future.

ACKNOWLEDGMENTS

We are grateful to Francesco Capparelli for advice on EU legislation and to Mario Raciti for advice on semantic relations. Gianpietro Castiglione acknowledges a studentship by Intrapresa S.r.l. and Italian “Ministero dell’Università e della Ricerca” (D.M. n. 352/2022).

REFERENCES

- [1] 2023. Repo tables. <https://anonymous.4open.science/r/nis-tables-CBB1>
- [2] 2023. Types of legislation. https://european-union.europa.eu/institutions-law-budget/law/types-legislation_en
- [3] Sam Arts, Jianan Hou, and Juan Carlos Gomez. 2021. Natural language processing to identify the creation and impact of new technologies in patent text: Code, data, and new measures. *Research Policy* 50, 2 (2021), 104144. <https://doi.org/10.1016/j.respol.2020.104144>
- [4] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *J. Mach. Learn. Res.* 3, null (mar 2003), 993–1022.
- [5] Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, Nikolaos Aletras, and Ion Androutsopoulos. 2020. LEGAL-BERT: The Muppets straight out of Law School. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics, Online, 2898–2904. <https://doi.org/10.18653/v1/2020.findings-emnlp.261>
- [6] EU Commission. 2023. NIS 2 Directive. <https://eur-lex.europa.eu/eli/dir/2022/2555/oj>
- [7] Spacy developers. 2023. spacy-clause library. <https://spacy.io/universe/project/spacy-clause>
- [8] Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. (2017). To appear.
- [9] Hamed Jelodar, Yongli Wang, Chi Yuan, Xia Feng, Xiahui Jiang, Yanchao Li, and Liang Zhao. 2018. Latent Dirichlet Allocation (LDA) and Topic modeling: models, applications, a survey. arXiv:1711.04305 [cs.IR]
- [10] Daniel Martin Katz, Michael J Bommarito II au2, and Josh Blackman. 2014. Predicting the Behavior of the Supreme Court of the United States: A General Approach. arXiv:1407.6333 [physics.soc-ph]
- [11] Jieh-Sheng Lee and Jieh Hsiang. 2019. Patent Claim Generation by Fine-Tuning OpenAI GPT-2. arXiv:1907.02052 [cs.CL]
- [12] Marco Lippi, Przemysław Pałka, Giuseppe Contissa, Francesca Lagioia, Hans-Wolfgang Micklitz, Giovanni Sartor, and Paolo Torroni. 2019. CLAUDETTE: an automated detector of potentially unfair clauses in online terms of service. *Artificial Intelligence and Law* 27, 2 (feb 2019), 117–139. <https://doi.org/10.1007/s10506-019-09243-2>
- [13] Pedro Henrique Luz de Araujo, Teófilo Emídio de Campos, Fabricio Ataiades Braz, and Nilton Correia da Silva. 2020. VICTOR: a Dataset for Brazilian Legal Documents Classification. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*. European Language Resources Association, Marseille, France, 1449–1458. <https://aclanthology.org/2020.lrec-1.181>
- [14] Masha Medvedeva, Michel Vols, and Martijn Wieling. 2020. Using machine learning to predict decisions of the European Court of Human Rights. *Artificial Intelligence and Law* 28, 2 (01 Jun 2020), 237–266. <https://doi.org/10.1007/s10506-019-09255-y>

- [15] Packtpub. [n. d.]. Packtpub object extraction. <https://subscription.packtpub.com/book/data/9781838987312/2/ch02lv1sec16/extracting-subjects-and-objects-of-the-sentence>
- [16] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [17] Octavia-Maria Sulea, Marcos Zampieri, Shervin Malmasi, Mihaela Vela, Liviu P. Dinu, and Josef van Genabith. 2017. Exploring the Use of Text Classification in the Legal Domain. arXiv:1710.09306 [cs.CL]
- [18] Fahad ul Hassan and Tuyen Le. 2020. Automated Requirements Identification from Construction Contract Documents Using Natural Language Processing. *Journal of Legal Affairs and Dispute Resolution in Engineering and Construction* 12, 2 (2020), 04520009. [https://doi.org/10.1061/\(ASCE\)LA.1943-4170.0000379](https://doi.org/10.1061/(ASCE)LA.1943-4170.0000379)
- [19] Nuo Xu, Pinghui Wang, Long Chen, Li Pan, Xiaoyan Wang, and Junzhou Zhao. 2020. Distinguish Confusing Law Articles for Legal Judgment Prediction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 3086–3095. <https://doi.org/10.18653/v1/2020.acl-main.280>
- [20] Haoxi Zhong, Chaojun Xiao, Cunchao Tu, Tianyang Zhang, Zhiyuan Liu, and Maosong Sun. 2019. JEC-QA: A Legal-Domain Question Answering Dataset. arXiv:1911.12011 [cs.CL]

APPENDIX

ARTICLE 11 OBJECTS

Item 11.1

Manual.

N11.1.1

- N11.1.2.a** a high level of availability of their communication channels by avoiding single points of failure, and shall have several means for being contacted and for contacting others at all times; they shall clearly specify the communication channels and make them known to constituency and cooperative partners
- N11.1.2.b** located at secure sites
- N11.1.2.c** with an appropriate system for managing and routing requests
- N11.1.2.d** the CSIRTs shall ensure the confidentiality and trustworthiness of their operations
- N11.1.2.e** availability of their services and their staff is trained appropriately
- N11.1.2.f** redundant systems and backup working space to ensure continuity of their service

ClausE.

S11.1.1

- S11.1.2.a** a high level of availability of their communication channels
- S11.1.2.b** NONE
- S11.1.2.c** NONE
- S11.1.2.d** the CSIRTs shall ensure the confidentiality and trustworthiness of their operations
- S11.1.2.e** availability of their services
- S11.1.2.f** NONE

Item 11.2

Manual.

- N11.2.1** that their CSIRTs jointly have the technical capabilities necessary to carry out the tasks referred to in paragraph 3

- N11.2.2** that sufficient resources are allocated to their CSIRTs to ensure adequate staffing levels for the purpose of enabling the CSIRTs to develop their technical capabilities

ClausE.

- S11.2.1** that their CSIRTs jointly have the technical capabilities necessary to carry out the tasks referred to in paragraph 3
- S11.2.2** that sufficient resources are allocated to their CSIRTs to ensure adequate staffing levels for the purpose of enabling the CSIRTs to develop their technical capabilities

...

ARTICLE 23 OBJECTS

Item 23.1

Manual.

- N23.1.1** that essential and important entities notify, without undue delay, its CSIRT or, where applicable, its competent authority in accordance with paragraph 4 of any incident that has a significant impact on the provision of their services as referred to in paragraph 3
- N23.1.2** the recipients of their services of significant incidents that are likely to adversely affect the provision of those services
- N23.1.3** that those entities report, inter alia, any information enabling the CSIRT or, where applicable, the competent authority to determine any cross-border impact of the incident

- N23.1.4** the notifying entity to increased liability

ClausE.

- S23.1.1** that essential and important entities notify, without undue delay, its CSIRT or, where applicable, its competent authority in accordance with paragraph 4 of any incident that has a significant impact on the provision of their services as referred to in paragraph 3
- S23.1.2** the recipients of their services of significant incidents that are likely to adversely affect the provision of those services
- S23.1.3** that those entities report, inter alia, any information enabling the CSIRT or, where applicable, the competent authority to determine any cross-border impact of the incident

- S23.1.4** the notifying entity to increased liability

...

Item 23.8

Manual.

- N23.8.1** notifications received pursuant to paragraph 1 to the single points of contact of other affected Member States.

ClausE.

- S23.8.1** NONE

Item 23.9

Manual.

- N23.9.1** a summary report, including anonymised and aggregated data on significant incidents, incidents, cyber threats and near misses notified in accordance with paragraph 1 of this Article and with Article 3
- N23.9.2** technical guidance on the parameters of the information to be included in the summary report
- N23.9.3** the Cooperation Group and the CSIRTs network about its findings on notifications received every six months

Clause.

- S23.9.1** a summary report, including anonymised and aggregated data on significant incidents, incidents, cyber threats and near misses notified in accordance with paragraph 1 of this Article and with Article 3
- S23.9.2** technical guidance on the parameters of the information to be included in the summary report
- S23.9.3** the Cooperation Group and the CSIRTs network about its findings on notifications received every six months

Item 23.10

Manual.

- N23.10.1** to the competent authorities under Directive (EU) 2022/2557 information about significant incidents, incidents, cyber threats and near misses notified in accordance with paragraph 1 of this Article and with Article 30 by entities identified as critical entities under Directive (EU) 2022/2557

Clause.

S23.10.1 NONE

Item 23.11

Manual.

- N23.11.1** implementing acts further specifying the type of information, the format and the procedure of a notification submitted pursuant to paragraph 1 of this Article and to Article 30 and of a communication submitted pursuant to paragraph 2 of this Article

Clause.

S23.11.1 NONE