A Methodology based on Commonsense Knowledge and Ontologies for the Automatic Classification of Legal Cases

Nicola Capuano
Department of Information and Electric Engineering and Applied Mathematics
University of Salerno
Fisciano (SA), Italy
ncapuano@unisa.it

Carmen De Maio
Department of Information and Electric Engineering and Applied Mathematics
University of Salerno
Fisciano (SA), Italy
cdemaio@unisa.it

Saverio Salerno
Department of Information and Electric Engineering and Applied Mathematics
University of Salerno
Fisciano (SA), Italy
salerno@unisa.it

Daniele Toti
Department of Information and Electric Engineering and Applied Mathematics
University of Salerno
Fisciano (SA), Italy
dtoti@unisa.it

ABSTRACT

We describe a methodology for the automatic classification of legal cases expressed in natural language, which relies on existing legal ontologies and a commonsense knowledge base. This methodology is founded on a process consisting of three phases: an enrichment of a given legal ontology by associating its terms with topics retrieved from the Wikipedia knowledge base; an extraction of relevant concepts from a given textual legal case; and a matching between the enriched ontological terms and the extracted concepts. Such a process has been successfully implemented in a corresponding tool that is part of a larger framework for self-litigation and legal support for the Italian law.

Categories and Subject Descriptors

I.2.7 [**Artificial Intelligence**]: Natural Language Processing—*Text analysis*; J.1 [**Computer Applications**]: Administrative Data Processing—*Law*

1. INTRODUCTION

Recent developments in artificial intelligence, in particular in the field of knowledge representation and reasoning, provide the support for legislative processes via the development of legal knowledge models. The successful exploitation of these models, combined with effective reasoning algorithms to derive information from them, can pave the way to bring about innovative systems able to assist users in legal matters and somewhat automatize long and often tedious "offline" activities.

This paper proposes a methodology to automatically classify descriptions of legal issues or experiences, expressed in natural language, by associating them with known legal ontologies. Specifically, this methodology is made up of a three-phase process, where: (i) a given ontology is firstly enriched by means of a "wikification" mechanism, that basically takes advantage of the Wikipedia commonsense knowledge base in order to expand the ontological terms with additional elements, or labels; (ii) the input text is "wikified" as well, so that relevant concepts or keywords that characterize it are identified and extracted from it; and (iii), the text is actually classified with respect to the given ontology, by matching the enriched ontological terms with the concepts extracted from the input text, thus returning a number of legal issues each with a corresponding score according to its relevance for the text.

The paper is organized as follows. In Section 2 related work is discussed. Section 3 describes the classification process, with its aforementioned three phases. Section 4 presents an application of the classification process as implemented in a tool which is part of a larger framework for self-litigation and legal support, while in Section 5 experimental results are reported. Finally, in Section 6 we draw our conclusions.

2. RELATED WORK

Over the years, the rising number of law disputes both in-court and out-of-court, as determined by recent studies carried out by the European Commission as well as by the Department of Justice of the United States of America, and the ever-increasing complexity of the legal domain itself, have determined in the Western world a surge in the practices of Alternative Dispute Resolution, and have consequently made the role of information systems more and more prominent. Even so, significant leaps forward are still hard to be achieved by the current systems available on the

WIMS'14 June 2-4, 2014 Thessaloniki, Greece

http://dx.doi.org/10.1145/2611040.2611048

market. In fact, commercial products usually only provide Internet-based tools for video conferencing, instant messaging and some templates for the manual definition of legal cases, whereas research initiatives are more focused towards the development of intelligent algorithms and solutions for aiding in the resolution of disputes. An early attempt to use computational intelligence approaches in this regard is represented by DEUS [21], which was meant to compute the agreement level in a family law property negotiation, while more sophisticated systems subsequently proposed include Split-Up [13] and Family-Winner [3], respectively using a rule-based mechanism combined with neural networks and a game theory-based approach. A more recent proposal is represented by the BEST project [15], leveraging on semantic web technologies like ontology-based search to support the retrieval of law cases. In the latest years, the use of such semantic technologies has indeed thrived, becoming more and more pivotal for implementing human-machine communication, and giving the concrete chance of fully harnessing the expressive power of ontologies. As a matter of fact, several ontologies have been defined for a wide range of knowledge domains (e.g. WordNet [20] for linguistics, Gene Ontology [14] for biology, just to name a relevant few), including law as well, being especially suited (and especially craving) for conceptual modeling, given the innermost complexity of the legal domain and the amount of data generated within its context. A number of significant legal ontologies have been thus described in literature, some even dating back to a time when semantic technologies were but in their embryo form. Among those, the Frame-based Ontology of Law (FBO) [17, 18] considers a legal system as fundamentally based on norms (either of conduct or of competence), acts and descriptions of concepts, and as such proposes a representation which takes into accounts these three types of elements. The FOLaw ontology, instead, employs a functional perspective, by relying on a structure made up of different kinds of "knowledge", ranging from normative knowledge, responsibility knowledge, up to reactive, creative, world and meta-level knowledge [16].[9] proposed a sort of "top ontology of the law", whose main assumption regarded law as a dynamic and an interconnected system of states of affairs, which evolved over time and where its constituents were strongly interconnected with one another. LRI-Core [4] tried to overcome the limits of FOLaw by including the latter's information and creating a structure made up of five "worlds": physical concepts (object and process), mental concepts, abstract concepts, roles and occurrences. A more recent Italian proposal is Jur-IWN [7, 8], which is basically an extension of the Italian version of WordNet in terms of a light-weight, lexical ontology specifically defined for the law domain. Even so, to the best of our knowledge there are few proposals in literature that tried to put together reasoning algorithms and legal ontologies in a cohesive solution for the automatic classification of legal cases, and even fewer that managed to succeed in the task; one of such attempts is for instance [2].

The methodology we propose exploits the information stored and modeled within a legal ontology and resorts to the commonsense knowledge drawn upon Wikipedia in order to enrich the ontological terms and extract concepts that characterize an unstructured text written in natural language. Such an approach uses a functionality based on Wikipedia, a so-called "wikification", which is able to derive additional information from an unstructured text by associating it with topics from the knowledge base, and thus link them to the text. This allows for either the augmentation of a given input text (for instance, an ontological term) or the identification of Wikipedia topics within it, whereas most of the methods for knowledge extraction usually employed in literature and in commercial products rely only on the information explicitly contained in the source text. In the following section we describe our classification process with greater detail.

3. THREE-PHASE CLASSIFICATION PROCESS

This section details the automatic classification process with its three phases building it up. Such a process takes as input a legal ontology, an input text (typically a paragraph describing a legal case or issue), and is meant to eventually produce a set of legal subjects treated in the given text, each with a corresponding score based upon the relevance of such a subject within the text itself. Figure 1 represents the whole classification process, which is composed of the following phases:

1. Ontology Enrichment, meant to extend the terms from a legal ontology with topics extracted from Wikipedia;
2. Concept Extraction, meant to identify and extract relevant concepts that characterize the input text;
3. Classification and Legal Issue Extraction, meant to match the enriched ontological terms with the extracted concepts from the text in order to return a list of legal issues relevant to the given text.

Further details of these phases are described in the following subsections.

3.1 Phase 1: Ontology Enrichment

The first phase of the classification process takes as input a legal ontology and proceeds with its wikification: each of the terms contained in it is "wikified" in order to add additional terms (or better, labels) to it, taken from Wikipedia, which are contextually related. We take advantage of the Wikipedia Miner tool [12] for performing the actual wikification. Table 1 shows the input and output of this phase.

Table 1: Input/Output of the Ontology Enrichment phase.

Phase 1: Ontology Enrichment	
Input	Legal ontology
Output	Legal ontology enriched via wikification, with additional terms added and scored according to their relevance

Specifically, for each term in the ontology, we select its label and description tags, if any, and we perform the wikification process upon them: for each of those terms, the wikification extracts a set of ⟨*topic, score*⟩ pairs, where *topic* corresponds to a Wikipedia article and *score* measures the relevance of such a topic with respect to the corresponding term. Let us clarify this via the following example, where we consider a term from a SKOS ontology with its description defined by a `skos:definition` element.

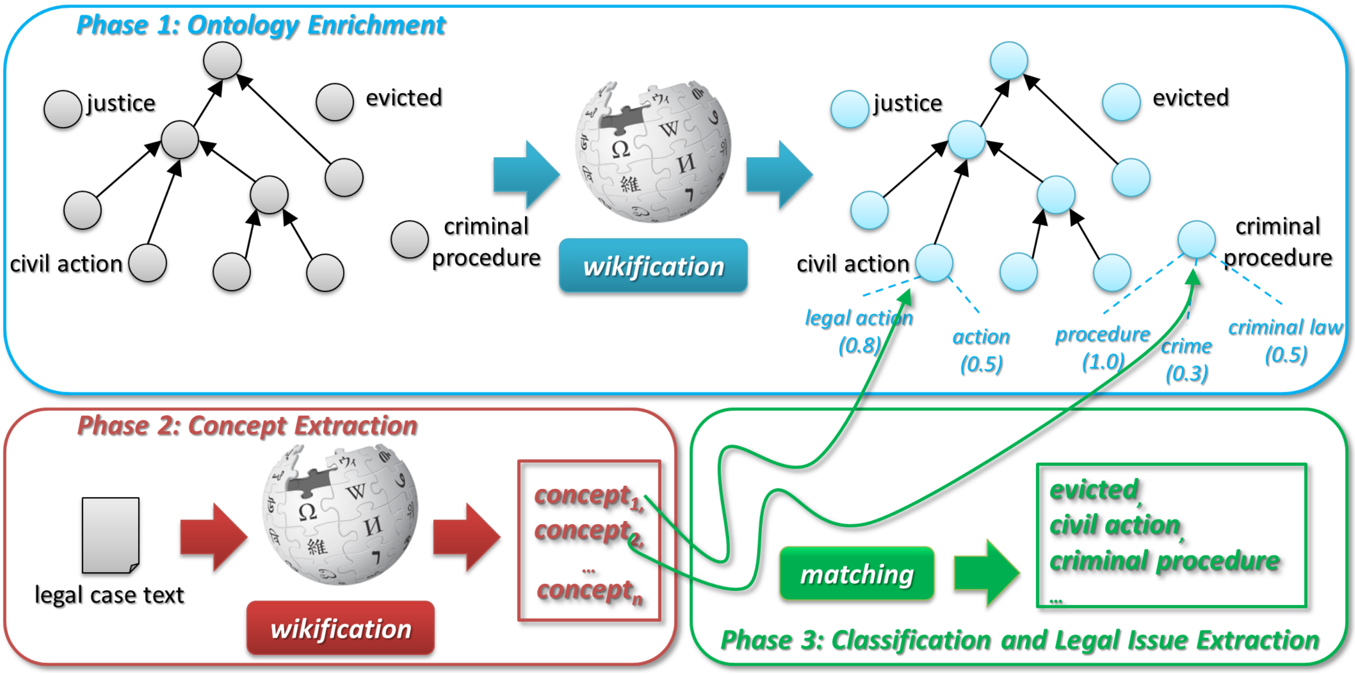

Figure 1: Schematization of the classification methodology seen as a three-phase process.

```
⟨rdf:Description rdf:about=
"http://eurovoc/namespace#VAT_Resource"⟩
    ⟨rdf:type rdf:resource=
    "http://www.w3.org/2004/02/skos/core#Concept"/⟩
    ⟨skos:broader rdf:resource=
    "http://eurovoc/namespace#own_resources"/⟩
    ⟨skos:definition xml:lang="it"⟩
     This comes from the application of a flat rate
     to the VAT base of each Member State.
     ⟨/skos:definition⟩
 ⟨/rdf:Description⟩
```

After performing the wikification process on this term, the extracted topics are added to it as additional labels, for the purpose of this example by using the `skos:hiddenLabel` elements. A threshold could be set upon the score of the returned topics in order to filter out those whose relevance is lower than a certain value. The resulting enriched term is shown below.

```
⟨rdf:Description rdf:about=
"http://eurovoc/namespace#VAT_Resource"⟩
    ⟨rdf:type rdf:resource=
    "http://www.w3.org/2004/02/skos/core#Concept"/⟩
    ⟨skos:broader rdf:resource=
    "http://eurovoc/namespace#own_resources"/⟩
    ⟨skos:definition xml:lang="it"⟩
     This comes from the application of a flat rate
     to the VAT base of each Member State.
     ⟨/skos:definition⟩
    ⟨skos:hiddenLabel⟩VAT resource%230.9
    ⟨/skos:hiddenLabel⟩
    ⟨skos:hiddenLabel⟩State%230.2
    ⟨/skos:hiddenLabel⟩
     ⟨skos:hiddenLabel⟩Tax%230.7
    ⟨/skos:hiddenLabel⟩
⟨/rdf:Description⟩
```

The enriched ontological terms are then stored and used in the final phase of the classification process, as described in Section 3.3.

3.2 Phase 2: Concept Extraction

The second phase of the classification process is independent of the previous phase and involves the extraction of concepts from an unstructured text written in natural language, which typically describes a legal case or issue. In order to do this, the input text is wikified as well, with the purpose of extracting a set of ⟨*topic, score*⟩ pairs corresponding to Wikipedia articles that are related to the text itself. This mechanism is in principle identical to the one used in Phase 1: the idea is to use the Wikipedia knowledge base to get a better insight on the given text. While in Phase 1 the topics returned by applying the wikification upon a term's label or description helped us enrich it with additional information, here we regard the topics returned as concepts that characterize the given text. Table 2 sums up the input and output of this second phase of the classification process.

Table 2: Input/Output of the Concept Extraction phase.

Phase 2: Concept Extraction	
Input	An unstructured text written in natural language
Output	A set of ⟨*topic, score*⟩ pairs

Let us report an example by considering the following legal issue described in natural language.

"In 2008 I got divorced and I became qualified for divorced benefits. In 2009 my ex-husband has remarried and then he died in 2011. Am I entitled to request allocation of a portion of the survivor's pension?"

The wikification process extracts from the above text a number of ⟨*topic, score*⟩ pairs, which as we said earlier do represent for our purposes the characterizing concepts of the input text, as seen in the following sample output.

```
⟨Wikification⟩
    ⟨Request input="In 2008 I got divorced and I became
    qualified for divorced benefits. In 2009 my ex-husband has
    remarried and then he died in 2011. Am I entitled
    to request allocation of a portion of the survivor's pension?
    If so, what percentage?" language="en"/⟩
    ⟨Response⟩
        ⟨WikifiedText⟩
            ⟨![CDATA[In 2008 I got [[divorce|divorced]] and I
            became qualified for divorced [[benefits]].
            In 2009 my ex-husband has remarried and then he
            died in 2011. Am I entitled to request allocation
            of a portion of the survivor's [[pension]]?
            If so, what percentage?]]⟩
        ⟨/WikifiedText⟩
        ⟨WikipediaTopics⟩
            ⟨WikipediaTopic id="2827896" title="Divorce"
            score="0.727"/⟩
            ⟨WikipediaTopic id="113578" title="Benefit"
            score="0.706"/⟩
            ⟨WikipediaTopic id="2835673" title="Pension"
            score="0.68"/⟩
        ⟨/WikipediaTopics⟩
    ⟨/Response⟩
⟨/Wikification⟩
```

3.3 Phase 3: Classification and Legal Issue Extraction

Once the chosen legal ontology has been enriched and the characterizing concepts have been extracted from the input text accordingly, as described respectively in Section 3.1 and 3.2, the third and last phase of the classification process takes care of actually classifying the input text on the enriched ontology, so that a number of legal issues or elements can be associated with it. Basically, this is obtained by carrying out a matching between the concepts characterizing the text and the terms of the enriched ontology. The input and output of this last phase are summarized in Table 3.

Table 3: Input/Output of the Classification and Legal Issue Extraction phase.

Phase 3: Classification and Legal Issue Extraction	
Input	- Legal ontology enriched; - A set of ⟨*topic, score*⟩ pairs characterizing the input text
Output	A number of legal issues from the enriched ontology relevant to the input text

The matching is performed by a weighted combination of the standard measures of Precision and Recall [19]. Typically, Precision and Recall are metrics used to measure the performance of an Information Retrieval System; nonetheless, in our case they are used to compute the matching, as described more formally below.

1. Let $A=\{a_1, a_2, \ldots, a_n\}$ the set of extracted concepts (i.e. Wikipedia topics), and $S=\{s_1, s_2, \ldots, s_n\}$the set of corresponding scores, resulting from the wikification process;

2. Let $C=\{c_1, c_2, \ldots, c_n\}$ the set of terms of the enriched ontology where $c_i = (\langle a^c_1, s^c_1 \rangle, \langle a^c_2, s^c_2 \rangle, \ldots, \langle a^c_s, s^c_s \rangle)$, corresponding to the set of pairs ⟨*topic, score*⟩ resulted from the enrichment of the term t_i;

3. Let $T = (\langle a^t_1, s^t_1 \rangle, \langle a^t_2, s^t_2 \rangle, \ldots, \langle a^t_m, s^t_m \rangle)$ the set of pairs ⟨*topic, score*⟩ output of the Concept Extraction phase from the input text T.

For each concept $c_i \in C$, we compute the Recall and Precision values as:

$$Precision = \frac{\sum(w^c_j * w^t_k)}{\sum w^c_j} \quad and \quad Recall = \frac{\sum(w^c_j * w^t_k)}{\sum w^t_k} \tag{1}$$

$\forall$ $(a^c_j, s^c_j) \in c_i$ and $\forall$ $(a^T_k, s^T_k) \in T$ so that $a^c_j = a^T_k$

Given that, let $\lambda_1 = \frac{1}{|A^t|}$ and $\lambda_2 = 1 - \frac{1}{|A^t|}$, where $|A^t|$ represents the number of concepts (topics) extracted from the input text, the matching is computed as:

$$\mu = \lambda_1 Recall + \lambda_2 Precision \tag{2}$$

As a result, we get a list of matching values μ, one for each enriched ontological term. Eventually, the final result includes all of those concepts whose value μ is greater than a specific threshold (which must be fine-tuned accordingly).

4. IMPLEMENTATION

The methodology described in Section 3 has been implemented in a prototype software tool that is part of a self-litigation framework developed within the eJRM project (described in [1] as well). Being the latter an Italian initiative, the tool has been applied to classify legal excerpts written in the Italian language upon a list of ontologies from different Italian and European sources: these include EuroVoc [6], a legal thesaurus, Italgiure [10], an Italian Law portal, and Dejure [5], an integrated online system for legal information. Since the commonsense knowledge base of Wikipedia is available in several languages, the tool has been successfully able to cope with the Italian language and produce corresponding results.

Figure 2 shows a screenshot of its current graphical user interface (GUI). As we can see from the uppermost part of the screenshot, the user can initially select one of the available ontologies and input the text in the box below. On the rightmost part of the GUI it is possible to browse the selected ontology in its entirety, while pressing the "Classify" button activates the classification: the results of the whole process are shown in the box in the middle in terms of matched ontological terms, each with its corresponding score, whereas the bottom box displays the concepts extracted from the text during Phase 2 with their scores as well.

5. EXPERIMENTAL RESULTS

In order to evaluate the effectiveness of the classification process described so far, we applied it on a subset of descriptions of legal cases manually classified by domain experts. Specifically, 100 descriptions of legal cases have been analyzed and classified with legal ontology terms, taken from Italgiure, the most complete ontology of the three mentioned above (EuroVoc, Italgiure and Dejure).

Table 4 wraps up the experimental results, where we show the legal cases or descriptions used for testing grouped by their length, with the corresponding average values of Precision, Recall and F-measure (the latter computed as $2 * ((P * R)/(P + R))$) for each group. Throughout our experimentation, we have considered legal cases and descriptions made up of short paragraphs whose length ranged from a few words to 1500 characters, as this was typically a sensible length for a user's input in this regard.

The overall effectiveness of the classification process has been then measured in terms of the micro-averaged Recall

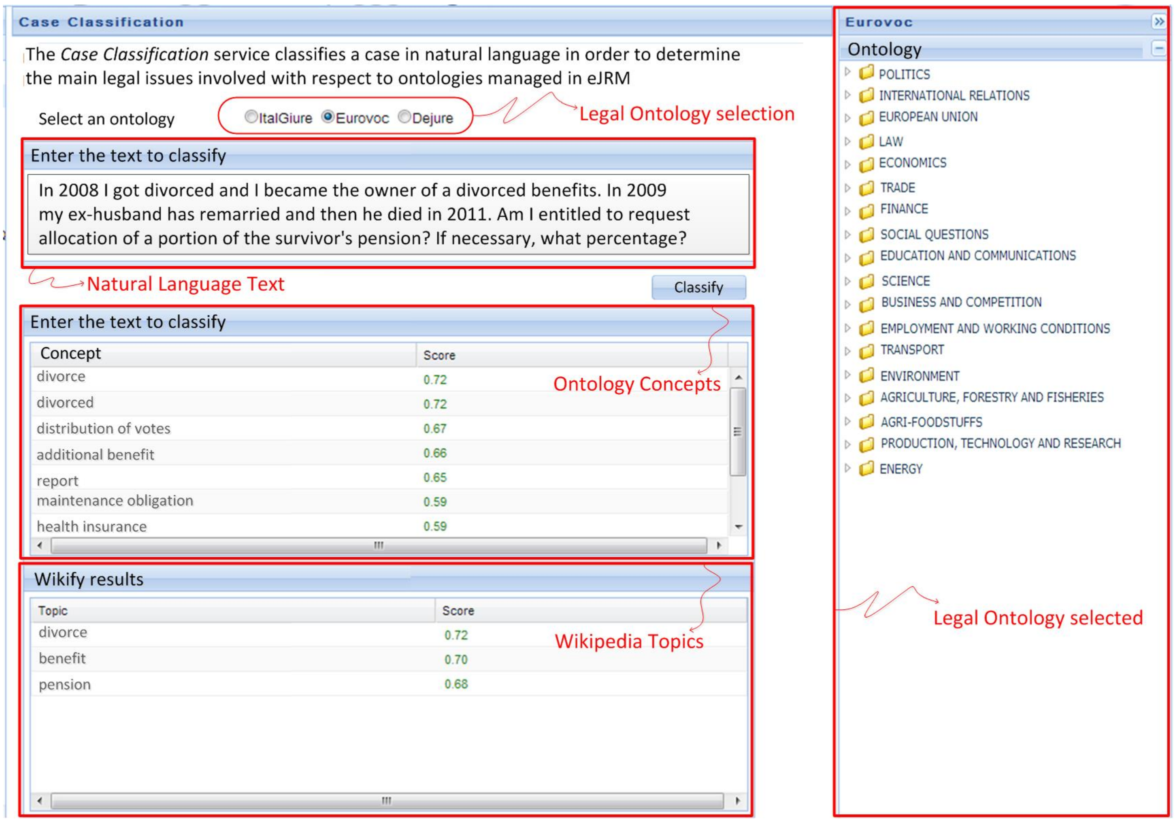

Figure 2: Screenshot of the software tool implementing the classification methodology.

Table 4: Results of the classification process grouped by the length of the test cases, with their related Precision (P), Recall (R) and F-measure (F) values.

# Texts	Length (char.)	P	R	F
15	[0, 300]	0.388	0.464	0.422
20	[301, 500]	0.546	0.574	0.559
25	[501, 700]	0.536	0.571	0.552
20	[701, 1000]	0.541	0.575	0.557
20	[1001, 1500]	0.632	0.642	0.636

and Precision [19]. Formally, let $D'= \{D_1, D_2, \ldots, D_n\}$ be a set of n descriptions of legal cases, O the set of terms that best classify them. For each description D_i, we consider $\lambda = 15$ steps up to its maximum recall value, and then measure the number of correct responses returned at each step λ. According to [19] the micro-average at the generic step λ is defined as:

$$Recall_\lambda \sum_{T_i} \frac{|R_{T_i} \bigcap B_{\lambda,T_i}|}{|R|} \quad Precision_\lambda \sum_{T_i} \frac{|R_{T_i} \bigcap B_{\lambda,T_i}|}{|B_\lambda|} \tag{3}$$

where R_{T_i} is the set of correct responses for a given description T_i, and B_λ the set of returned responses at the step λ, for the test T_i.

It is important to stress out that the effectiveness and accuracy of the classification process, as it currently stands, is strongly dependent on the quality of the underlying legal ontology selected, along with the wikification mechanism. The latter, as testified by the results, improves its effectiveness as the length of the input text increases, since the more information is available, the better the tool is able to correctly understand its context and come up with meaningful concepts. Besides, given the vastness and breadth of the Wikipedia knowledge base, while on one hand recall is somewhat easier to achieve, some loss of precision might be detected when dealing with polysemic words (both in the input text and among the ontological terms) that possess different meanings according to the specific context: that is why short sentences are harder to be correctly wikified, whereas longer paragraphs tend to produce increasingly better results. In the presence of more technical, less ambiguous terms, instead, the classification process treads an easier road and has less trouble in identifying the correct concepts accordingly. Nevertheless, because the extracted concepts are eventually matched with those appearing in the legal ontology, the matching phase (and consequently the use of a "controlled dictionary" against which such concepts are searched) mostly makes up for the potential loss in precision of the wikification mechanism.

Performance in terms of execution time is almost instantaneous for texts of the considered length, and stays this way unless the length of the input text grows exceedingly larger (e.g. long pages of text), which is nonetheless an unrealistic event given the nature of the legal descriptions taken into account; for longer legal cases, additional tests will have to be executed and splitting mechanisms could be employed should the need arise.

6. CONCLUSION

In this paper we have proposed a methodology for the automatic classification of legal cases written in natural language, given the availability of a legal ontology and via a three-phase process. Its core lies in resorting to a wikification mechanism based on the Wikipedia knowledge base, which on one hand enables us to enrich the terms of an existing ontology, and on the other hand allows for the identification of characterizing concepts from an unstructured text; a matching between the enriched ontological terms and those concepts finally returns the legal issues related to the input text, all in an automatic way. This methodology in principle supports all the languages of Wikipedia, and has been implemented and tested within a framework for legal support for the Italian law.

Results are encouraging, since no human intervention is involved in the classification process: the combined metrics of Precision and Recall reach and exceed values of 0.6 with legal cases described by more than 1000 characters. At the same time, this kind of approach suffers from the same intrinsic limits of other automatic processes, and as such does not strive to completely replace the know-how of a legal domain expert; instead, it is meant to complement and enhance the work of both users and legal professionals, by providing them with a tool able to speed up and support their respective legal matters.

Improvements to the proposed methodology and its current implementation may include a refinement of the wikification process by potentially establishing specific linguistic and domain-based rules for fine-tuning the classification for very specific legal areas, in accordance with a corresponding well-defined ontology. Furthermore, experimentation on longer texts is needed and is currently underway, by which it will be possible to both assess the extent to which precision and recall rise (since they tend to increase with the length of the text) and when (or if) they somehow flex if the considered text grows too much, so that additional refinements might be applied in order to further strengthen the results achieved so far.

Acknowledgments

This work has been partially supported by the eJRM ("electronic Justice Relationship Management") project, co-funded by the Italian Ministry of Instruction and Research (ref.: PON01_01286).

7. REFERENCES

[1] Arosio, G., Bagnara, G., Capuano, N., Fersini, E., and Toti, D.: Ontology-driven Data Acquisition: Intelligent Support to Legal ODR Systems. In *International Conference on Legal Knowledge and Information Systems (JURIX)*, 2013.

[2] Bartolini, R., et al.: Automatic Classification and Analysis of Provisions in Italian Legal Texts: A Case Study. In *OTM Workshops 2004, LNCS 3292, pages. 593-604*, 2004.

[3] Bellucci, E. and Zeleznikow, J.: Family winner: A computerised negotiation support system which advises upon australian family law. In *ISDSS2001*, 2001, pages 74-85.

[4] Breuker, J.: Constructing a legal core ontology: Lri-core. *In Proceedings of the Workshop on Ontologies and their Applications, Sao Luis, Maranhao, Brazil*, 2004.

[5] Dejure. http://dejure.giuffre.it

[6] Eurovoc. http://eurovoc.europa.eu

[7] Gangemi, A., Prisco, A., Sagri, M.T., Steve, G., and Tiscornia, D.: Some ontological tools to support legal regulatory compliance, with a case study. In *Meersman, R. and Tari, Z., editors, On The Move to Meaningful Internet Systems 2003: OTM 2003 Workshops, volume 2889 of Lecture Notes in Computer Science, pages 607-620. Springer-Verlag*, 2003.

[8] Gangemi, A., Sagri, M.T., and Tiscornia, D.:Metadata for content description in legal information. *In 14th International Workshop on Database and Expert Systems Applications (DEXA'03), IEEE Computer Society*, 2003.

[9] Hage, J. and Verheij, B,: The law as a dynamic interconnected system of states of affairs: A legal top ontology. In *International Journal of Human-Computer Studies, 51(6):1043-1077*, 1999.

[10] Italgiure Web. http://www.italgiure.giustizia.it/

[11] Mihalcea, R., and Csomai, A.: Wikify!: linking documents to encyclopedic knowledge. In *Proceedings of the 16th ACM Conference on Information and Knowledge Management (CIKM)*, 2001.

[12] Milne, D. and Witten, I.H.: An open-source toolkit for mining Wikipedia. In *Proc. New Zealand Computer Science Research Student Conf., NZCSRSC*, 2011. http://wikipedia-miner.cms.waikato.ac.nz

[13] Stranieri, A. and Zeleznikow, J.: The Split_Up system: Integrating neural networks and rule-based reasoning in the legal domain. In *proc. of the Fifth International Conference on Artificial Intelligence and Law*, 1995.

[14] The Gene Ontology. http://www.geneontology.org

[15] Uijttenbroek, E.M., Lodder, A.R., Klein, M.C.A., Wildeboer, G.R., van Steenbergen, W.R., Sie, L.L., Huygen, P.E.M. and van Harmelen, F.: Retrieval of Case Law to Provide Layman with Information about Liability: Preliminary Results of the BEST-project. In *Computable Models of the Law*, 2008.

[16] Valente, A. and Breuker, J.: A functional ontology of law. In *Bargellini, G. and Binazzi, S., editors, Towards a Global Expert System in Law. Cedam Publishers, Padua, Italy*, 1994.

[17] van Kralingen, R. W.: Frame-based Conceptual Models of Statute Law. In *Computer/Law Series, 16, Kluwer Law International, The Hague*, 1995.

[18] van Kralingen, R. W.: A conceptual frame-based ontology for the law. In *Visser, P. and Winkels, R., editors, Proceedings of the First International Workshop on Legal Ontologies (LEGONT'97), pages 15-22*, 1997.

[19] van Rijsbergen, C.J.: Information Retrieval, second ed., Dept. of Computer Science, University of Glasgow, 1979.

[20] WordNet. http://wordnet.princeton.edu

[21] Zeleznikow, J., Meersman, R., Hunter, D. and van Helvoort, E.: Computer tools for aiding legal negotiation. In *proc. of the 6th Australasian Conference on Information Systems*, 1995.