Deliverable 4.2

Demo Version of the Pilot Tool

Grant Agreement nº:	101007420
Project Acronym:	ADELE
Project Title:	Analytics for DEcision of LEgal cases
Website:	https://site.unibo.it/adele
Contractual delivery date:	30/09/2022
Actual delivery data:	07/10/2022
Contributing WP	WP4
Dissemination level:	Public
Deliverable leader:	APIS EUROPA
Contributors:	UNIBO, UNITO

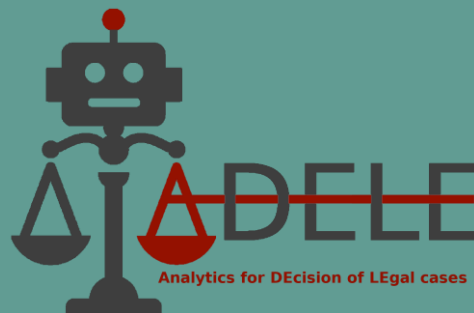**This project is funded by the European Union's Justice Programme (2014-2020)**

1. Document History

Version	Date	Author	Partner	Description
1.0	30/09/2022	Hristo Konstantinov	APIS	First draft
1.1	03/10/2022	Dilyana Bozhanova, Lilia Kachoreva	APIS	Review and minor revisions
1.2	04/10/2022	Hristo Konstantinov	APIS	Final draft
2.0	05/10/2022	Federico Galli	UNIBO	Review and addition of UNIBO task
2.1	06/10/2022	Galileo Sartor	UNITO	Addition of UNITO task
3.0	07/10/2022	Giovanni Sartor, Giuseppe Contissa, Federico Galli	UNIBO	Final version

2. Contributors

Partner	Name	Role	Contribution
APIS	Hristo Hristov, Daniel Traykov, Vasil Oreshenski, Julian Lazarov	WP leaders	Development of the demo version of the pilot tool
UNIBO (DiSi)	Paolo Torroni, Elena Palmieri, Giulia Grundler, Andrea Galassi	WP Participant	Development of models for automated argument extraction and outcome prediction
UNITO	Galileo Sartor, Luigi Di Caro	WP Participant	Network and citation analysis
UNIBO (Cirsfid)	Piera Santin Federico Galli, Alessia Fidelangeli, Giuseppe Contissa, Francesca Lagioia	WP Participant	Network and citation analysis; Corpus annotation and guidelines revisions.

Disclaimer: On behalf of APIS EUROPA, the present deliverable has been drafted by Hristo Konstantinov, Dilyana Bozhanova and Lilia Kachoreva.
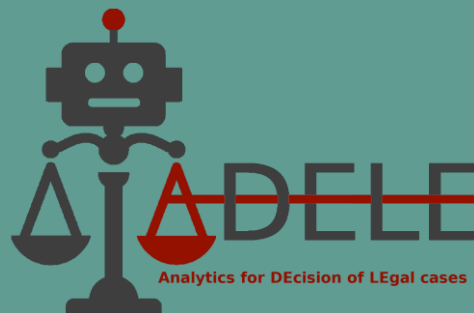
3. Table of Contents

4. List of Figures

5. List of Tables

6. List of Acronyms

AI	Artificial Intelligence
APIS	Apis Europa, beneficiary in ADELE project
CJEU	Court of Justice of the European Union
EU	European Union
EUI	European University Institute, beneficiary in ADELE project
LIBRe	LIBRe Foundation, beneficiary in ADELE project
ML	Machine Learning
NLP	Natural Language Processing
TM&P	Trademark and patents
UNIBO	University of Bologna, coordinator of ADELE project
VAT	Value Added Tax
WP	Work package
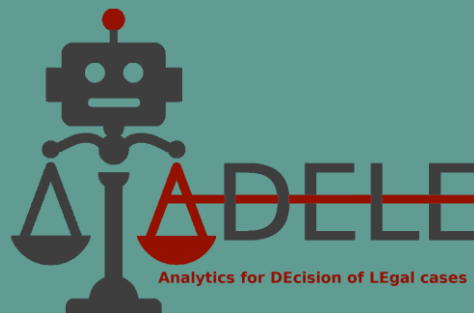
7. Introduction

This deliverable presents a demo version of the ADELE pilot tool, which is implemented as a web application freely accessible on https://adele-tool.eu.

The demonstrator incorporates the key content and features of the tool. The aim of this early version is to allow ADELE partners to demonstrate the provisional functionalities of the platform at two preliminary test events with judges in Italy and Bulgaria and gain feedback on its features. The received opinions and recommendations from the target users will facilitate and favour the development of the fully functional final version of the tool.

Since the pilot tool is a constantly evolving and updating legal database and software, the goal of the present document is to provide a brief overview on the status of the platform development at the stage of launching its demo version.

8. Brief Description of the Published Legal Content

8.1. Legislation

ADELE pilot tool provides access to the main instruments of EU law and the national legislations of Italy and Bulgaria in the chosen legal domains – VAT and Trademark and patents. Besides the consolidated texts of the acts currently in force, the tool's database includes a number of repealed legislative instruments that are often cited in judicial decisions.

Table 8-1 below provides information about the number of legislative instruments included in the demo version of the tool per legal domain and jurisdiction. Clicking the links in the table will open a list of documents in the pilot tool for the respective legal domain and jurisdiction.

	Value Added Tax	Trademark and patents
European Union	25	20
Bulgaria	16	10
Italy	9	8

Table 8-1. Number of legislative instruments included in the demo version of ADELE pilot tool

The demo version does not include time versions of the legislative instruments.
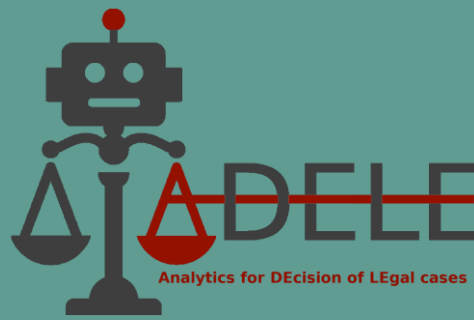
8.2. Case Law

The demo version of the pilot tool contains 861 anonymised and annotated decisions of Italian and Bulgarian courts as well as 263 non-annotated decisions of the Court of Justice of the European Union (CJEU) in the chosen legal domains – VAT and Trademark and patents. The annotations have been performed by legal experts of the partners UNIBO, EUI, APIS and LIBRe in XML format according to the annotation guidelines elaborated with Deliverable 2.1, later updated and finalised with Deliverable 2.2.

Table 8-2 below summarises the number of decisions included in the demo version of ADELE pilot tool per legal domain and jurisdiction. Clicking the links in the table will open a list of documents in the pilot tool for the respective legal domain and jurisdiction.

	Value Added Tax	Trademark and patents
Court of Justice of the EU	245	28
Bulgarian courts	217	179
Italian courts	235	230

Table 8-2. Number of court decisions included in the demo version of ADELE pilot tool

9. Current Status of the Implemented Functionalities

9.1. Legal databases functionalities

The ADELE pilot tool offers a number of basic functionalities for browsing and searching documents that a typical for legal databases.

1) Browsing documents

The ADELE platform provides access to four document collections divided by country and legal domain:

- Bulgaria – Value added tax, Trademark and patents
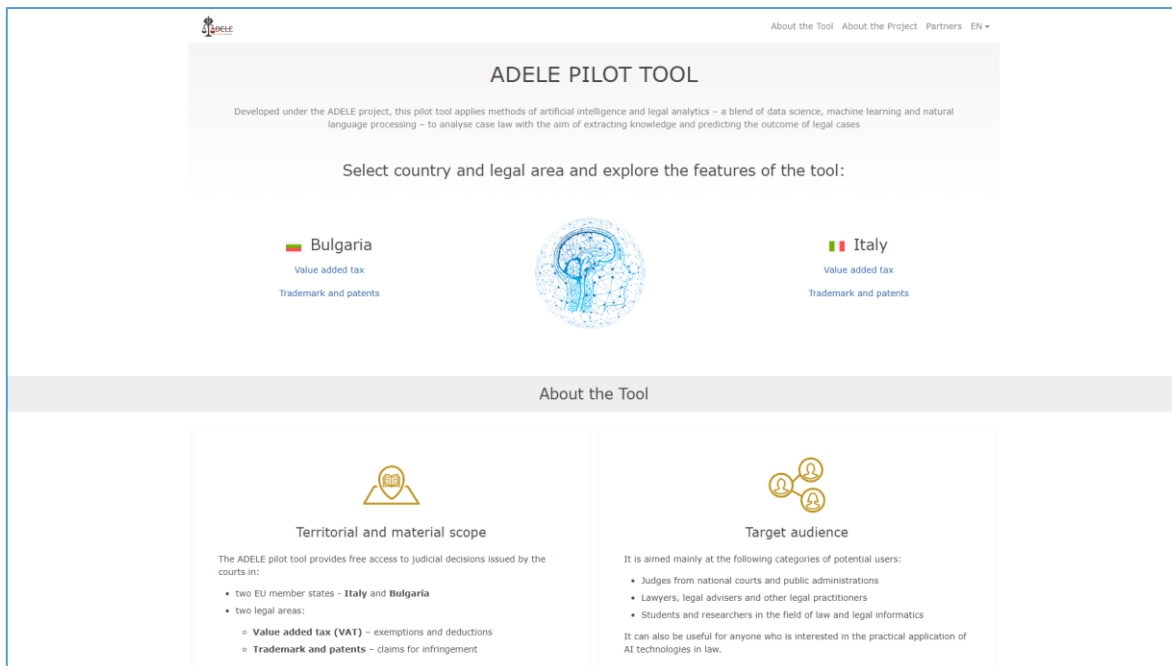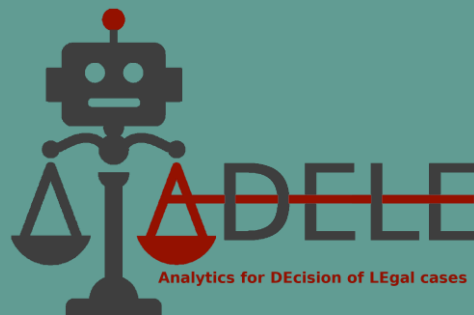- Italy – Value added tax, Trademark and patents.

Figure 9-1. Home screen of ADELE pilot tool with options for choosing country and legal domain

The choice of a particular document collection leads users to the main page of the tool where they can browse several lists of documents according to their type by clicking on any of the four buttons above the list:

- *National case law* – decisions of courts or other judicial authorities of Italy or Bulgaria in the chosen legal domain

- *National legislation* – main legislative instruments of Italy or Bulgaria in the chosen legal domain

- *EU case law* – relevant judgments of the CJEU in the chosen legal domain in Italian or Bulgarian language (depending on the selected country)

- *EU legislation* – relevant instruments of EU law in the chosen legal domain in Italian or Bulgarian language (depending on the selected country).

Documents in each list can be filtered by one or more criteria that are specific for the respective document collection, e.g. by court, case outcome or year of the decision.
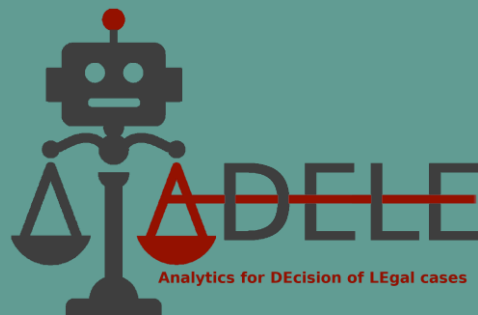
Figure 9-2. Main page of ADELE tool with lists of documents and filters

2) Searching documents and filtering search results

Users can perform searches in the four document collections by entering a keyword or a key phrase in the search box above the document list. The result is a list of documents containing the search expression. The list could be further filtered by using one or several filters.

Figure 9-3. Search results in ADELE tool
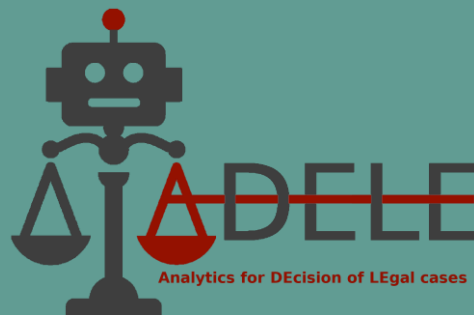
Opening a document in the list of search results displays the highlighted matches of the searched keywords within the text. The three buttons on the left side of the search box allow users to select next or previous match or remove the highlighting.

3) Table of contents

Where documents, such as instruments of national or EU legislation, have a complex hierarchical structure, users can navigate between the individual sections of the document via a table of contents in the panel on the left. When selecting an element in the table, the text of the document is positioned on the corresponding section. In addition, users can make quick searches in the text of the structural elements in the table of contents.

Figure 9-4. Table of contents in ADELE tool

4) Inline links to cited documents

Citations to other legal documents within the text of judicial decisions are provided in the form of inline links. Clicking on a link opens the cited document. If a particular provision of a legislative act is cited, then the text of the open document is positioned on the cited provision.
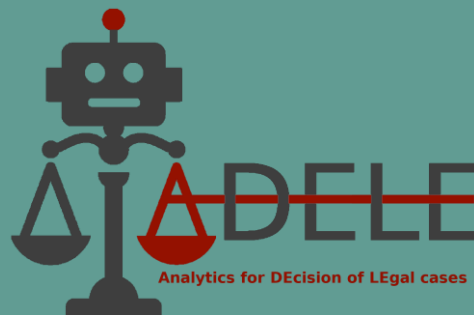
As part of the performed citation analysis, a special icon is placed next to each link. Clicking on this icon brings up a list of all documents that cite the same document or provision. This feature is especially convenient as it allows users to easily and quickly find other court decisions citing the same legal rule or judicial precedent.

9.2. AI and ML-powered functionalities

The demo version of ADELE pilot tool offers the following features that are based on the application of specific AI and ML methods and techniques:
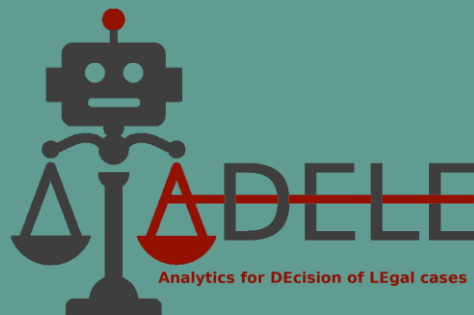
1) *Automatically extracted key terms*

Providing a set of relevant key terms (keywords) facilitates legal information retrieval by assisting users quickly to understand the main factual and legal issues discussed in a particular case without reading its text. In the demo version of the pilot tool, key terms are displayed under the title of judicial decisions in lists of documents as well as in an open document. Depending on the length of the decision, between 5 and 15 key terms are usually shown.

Figure 9-5. Automatically extracted key terms in a list of documents

Users can click on each key term in the list of the automatically extracted keywords, which leads to highlighting all its occurrences within the text of the decision.

The automated key term extraction is performed by a trained spaCy model based on Named Entity Recognition (NER). For training of the dataset two specialised dictionaries were used containing legal terms in the relevant fields of law – one in the field of VAT and one in the field of TM&P law. They were compiled by legal experts of APIS project team. In addition, a few more general law-related dictionaries such as EuroVoc are exploited. The applied model labels the key terms recognised in the text of a decision in two categories: the terms from the specialised dictionaries are labelled as priority key terms, whereas those contained in other dictionaries are labelled as non-priority key terms. The terms that the trained model classifies as priority key terms are directly included in the final selection, whereas the KeyBERT library (Grootendorst, 2020) is used for the non-priority keywords to extract only those keywords that are most relevant to the processed document. However, before using KeyBERT, additional filtering is applied for the non-priority key terms. For instance, only non-priority key terms that have at least two occurrences within the text of the processed document are included.

2) *Automatically extracted case summaries*

Similar to automatically extracted key terms, the aim of this feature is to help users quickly grasp the essence of the case by extracting key arguments (sentences or paragraphs) from the argumentative part of the decision. A summary is not generated for decisions that contain already a summary (an abstract) produced by a legal expert. During the experiments, it became obvious that it does not make much sense to generate a summary also in cases where the argumentative part of a decision is very short.
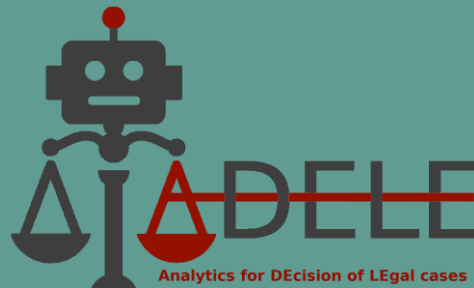
In the demo version of the pilot tool, the automatically generated summaries are displayed in lists of documents or in open documents under the title of judicial decisions, just below the automatically extracted key terms. Clicking on a sentence (or group of sentences) in a summary highlights this item within the text of the decision so that users can explore the context in which it was used.

The first experiments with text summarisation started by using the TextRank model[1], but the results were not satisfactory. Therefore, it was decided to train a spaCy model[2] on a dataset of about 40.000 Bulgarian court decisions summarised manually by legal experts. The model is based on sentence text categorisation, where the aforementioned manually annotated data were used to generate a

---

[1] https://cran.r-project.org/web/packages/textrank/vignettes/textrank.html.
[2] https://spacy.io/usage/models.

training dataset containing two categories of sentences: *summary* and *not_summary*. The applied approach led to a significant improvement of the results for the Bulgarian decisions.

For Italian decisions, initially we also tried to apply the TextRank model. Since again the results were not satisfactory, we used the trained "Bulgarian" spaCy model, because it was based on a multilingual transformer. In this way, we received far better results.

Figure 9-6. Automatically extracted case summaries in a list of documents

3) Citation analysis

The demo version of ADELE pilot tool provides an initial implementation of several features based on the citation and network analysis performed by the ADELE partners.

First of all, the recognised legal citations to EU and national legislation and case law in the text of the decisions appear in the form of in-line links. Clicking such a link opens the text of the cited document either in ADELE tool (when it is available there in full text) or in another online legal database (e.g., EUR-Lex portal).

Secondly, when clicking on the small icon next to an on-line link, a new browser window opens containing a list of documents that cite the same judicial decision or legislative act, respectively the
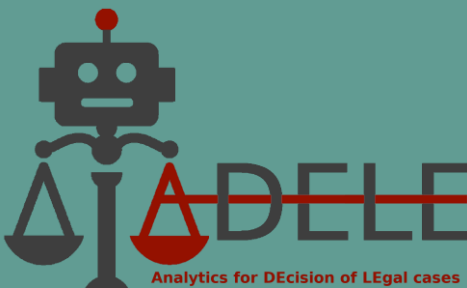
same provision of the legislative act. Thus, users are able to find potentially comparable cases that refer to the same legal document and discuss similar legal issues.

Thirdly, based on in- and out-citations, for each open document users can produce lists of all documents that are citing or are cited by that particular document. To this end, they could use the two buttons above the text of the open document: "Links to the document" or "Links from the document".

Figure 9-7. Examples of in-line links and the buttons for in- and out-citations

Finally, the demo version of ADELE pilot tool provides an option for the user to assess the importance of national and EU cases by receiving statistics of the most cited documents. This feature is available by clicking the "Citation Analysis" button on the top of the screen. It produces two lists of the most cited decisions of national courts and judgments of the CJEU for the respective legal domain and country (e.g. VAT decisions of Bulgarian courts) ranked by citation frequency. The number in parentheses before the title of a decision indicates the number of documents in the ADELE database that cite that particular decision. Clicking on the small icon before the title of the decision generates a list of all decisions citing it.

The citations were initially extracted by using regular expressions to capture the different formulations of citations in the analysed legal documents, with the assistance of legal experts. Once
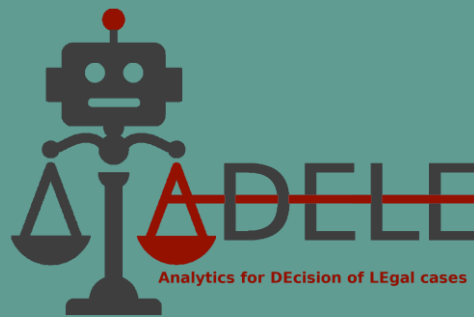
extracted, metadata regarding these cases was searched in different publicly available judicial databases. This information was collected in a database, keeping track of the citations.

Once the collection of citations was complete, the data was analysed with different network analysis tools. Initially, we were interested in sorting the cases by number of citations (direct and indirect). To do so, we found that the *degree centrality* algorithm was the most effective.

The network of citations will be further enhanced by adding more decisions based on the expansion of the dataset. Furthermore we will evaluate its usefulness in combination with the ontological framework, which is currently being developed.

Figure 9-8. Citation analysis: list of most cited Bulgarian VAT decisions and judgments of the CJEU

4) *Visual representation and highlighting of decision annotations*

As already pointed out in Section 8.2 above, the demo version of ADELE pilot tool contains the full set of XML files of Italian and Bulgarian decisions manually annotated by legal experts of the partners UNIBO, EUI, APIS and LIBRe according to the annotation guidelines elaborated with Deliverables 2.1 and 2.2. These annotated data are used as training dataset for all AI and ML tasks performed within the project framework.
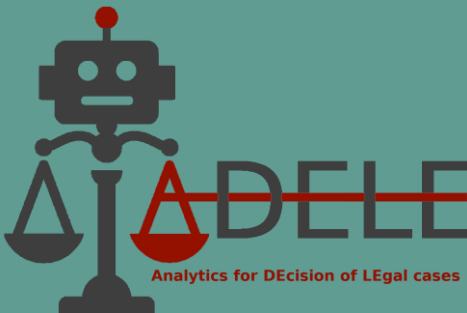
The demo version provides a special functionality for a structured visual representation of the annotated elements of the decisions allowing users to navigate between different sections and pieces of text by highlighting judicial arguments and other annotated units. The annotations are visualised in the panel on the left of the decision text. They are presented in a structured form similar to a table of contents. In addition, judicial arguments are classified by type (legal, factual or legal-factual premises) and argumentation scheme (e.g. argument from rule, from precedent, from interpretation, etc.). The panel on the left displays further different argument chains where one or more arguments is/are a premise of another, thus forming linked lines of argumentation.

Clicking on an annotated element in the panel on the left leads to highlighting of the annotated piece of text on the right. Thus, users can browse the annotated elements, identify them within the text and perceive contextually their position and importance. In particular, the arguments highlighting feature aims to assist users in perceiving and exploring visually the complex argument chains of the judicial reasoning.

Figure 9-9. Structured visual representation and highlighting of annotated elements

A small horizontal panel under the decision's text indicates interdependences between the highlighted element and other annotated elements of the same decision. For instance, for each selected claim users can visualise information on which request of what party is supported by that
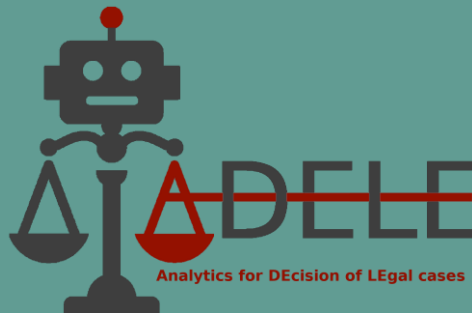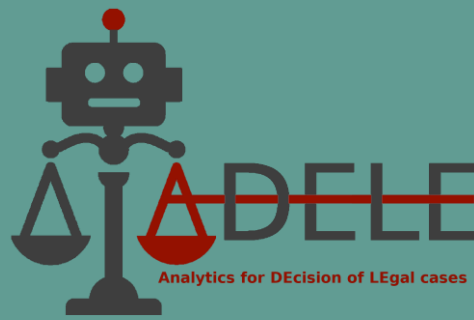
claim, which are the arguments of this party that support the claim and which findings of the court are related to that claim. As for judicial arguments, the below panel indicates the interconnections and interdependences between arguments (whether an argument supports/attacks or depends on another one).

5) *Automated argument extraction*

This feature makes it possible for users to copy-paste in a specially designed form the text of a judicial decision and ask the tool to automatically identify, highlight and classify judicial arguments.

Figure 9-10. Example for using the automated argument extraction functionality: input and output

The visual representation, highlighting and classification of the arguments (premises and conclusions) is identical with the same features implemented for the manual expert annotations of the decisions as already described in pt. 4) above.

In the demo version of the pilot tool, the automated argument extraction feature is available only for Italian VAT cases. In order to use it, users have to click on the "Argument Extraction" button on the top of the screen. For Bulgarian decisions and Italian TM&P cases, this functionality is made available as a mock-up only, since the developed model is still being tested.

The automated argument extraction functionality is based on the data annotated following the argument mining guidelines included in D2.1 and D2.2. Based on these data, four tasks were performed:

- Argument Detection: given a sentence, classify it as premise, conclusion, or neither;

- Argument Classification: given a sentence that is known to be argumentative, classify it as premise or conclusion;

- Type Classification: a multi-label classification problem where a sentence that is known to be a premise is classified as legal and/or factual;

- Scheme Classification: a multi-label classification task where a sentence, known to be a legal premise, is classified according to its argumentative scheme.

For all tasks we adopted three different representations of the input text, such as TF-IDF, Sentence-BERT (SBERT) (Reimers and Gurevych, 2019), Legal-BERT (Chalkidis et al., 2020). As classifiers, we have chosen a set of traditional machine learning techniques that have low computational requirements (e.g., Linear SVC, Random Forest, K Neighbors). The table below presents the average performance for each task considering both the embedding techniques and the classifier adopted.

Task	Embedding	Classifier	Score
Argument detection	TD-IDF	Linear SVC	0.70
Argument classification	TD-IDF	Random Forest	0.88
Type classification	LegalBert	SVC	0.85
Scheme classification	TD-IDF	Linear SVC	0.75

Table 9-1. Experimental results in the argument mining task

The tables containing complete results are presented and discussed in Grundler et al. (2022), accepted at ArgMining 2022: 9th Worhshop on Argument Mining.

In future work, we aim to explore the use of over-sampling and data augmentation techniques to overcome the detected unbalance of classes in each task. We also want to study the impact of pre-processing and the use of alternative classifiers such as logistic regression. Finally, we aim to
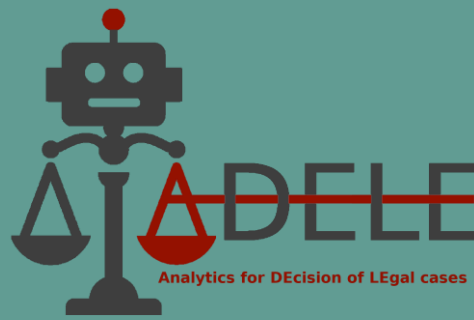
improve the robustness of our experimental findings, for example, by considering multiple seed runs or applying the method proposed by Lai et al. (2021).

6) Outcome prediction (mock-up)

As this feature is still in an experimental phase with Italian VAT cases, the partners decided to implement it in the demo version of the pilot tool as a mock-up only. The mock-up illustrates the input required from the user and the expected output in the form of a prediction of the outcome related to a particular request of party, possibly supported by particular claims and arguments. In order to use the mock-up, users have to click on the "Outcome Prediction" button on the top of the screen. Then, they have to type or copy-paste the request of a party to a particular legal case and one or more claims supporting this request in the text boxes of the "Outcome Prediction Form" that appears in a new window of the browser. Each claim, in turn, can be supported by one or more arguments. Claims and arguments can be added or removed by using the "Add claim / argument" and "Remove claim / argument" buttons. After typing or pasting the required input data concerning a request of a party to a legal case, users can click the "Predict" button in order to receive a prediction of the case outcome ("Uphold" or "Reject") in relation to this particular request and a related confidence score.

Figure 9-11. Example for using the outcome prediction functionality: input and output
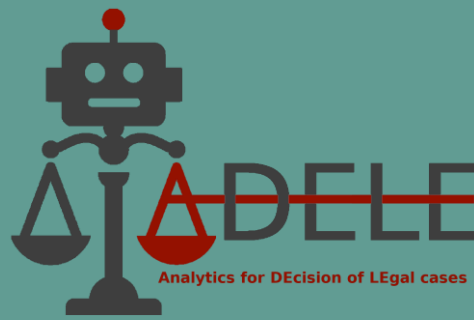
The outcome prediction functionality is based on training data annotated following annotation guidelines contained in D2.1 and D2.2. The ML task was to predict the court's decisions on a single party's request based on the information provided by the parties before the case. In particular, four experimental settings were considered: (1) request, claims, arguments; (2) request, claims, arguments, motivation; (3) request, claims, arguments, decision; (4) request, claims, arguments, motivation, decision. For each setting, we decided to perform two different experiments: (1) one in which we trained and test only on second-instance decisions, and (2) one where we also included first-instance decisions.

For each task, two representations techniques for input text were used, namely TF-IDF vectorization and Sentence-BERT (SBERT). As classifiers, we have chosen a set of traditional machine learning models with low computational requirements: Linear SVC, SVC, Random Forest, Gaussian Naive Bayes and K-Neighbours. The table below presents the average performance for each experimental setting considering both the embedding techniques and the classifier adopted.

Setting	w/o first-instance requests			w/ first-instance requests		
	Embedding	Classifier	Score	Embedding	Classifier	Score
Req, claims, args	SBERT	Linear SVC	0.68	TF-IDF	Linear SVC	0.69
Req, claims, args, mots	SBERT	Linear SVC	0.66	SBERT	Linear SVC	0.66
Req, claims, args, dec	SBERT	Linear SVC	0.68	TF-IDF	Random Forest	0.72
Req, claims, args, mots, dec	SBERT	Linear SVC	0.66	TF-IDF SBERT	Linear SVC	0.64

Table 9-2. Experimental results in the outcome prediction task

The tables containing complete results is presented and discussed in a paper recently submitted at JURIX 2022 Conference.

The ideal outcome of the outcome prediction task would consist in being able to predict the decision of the court on the basis of the information provided by the parties before the case. This task is approximated in our first experiments by predicting the outcome based on corresponding fragments in the narrative provided by courts, which we captured through the requests, claims, and arguments marked elements. To achieve this goal, we will experiment with more advanced techniques, such as
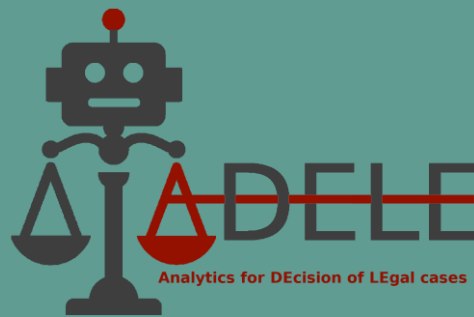
neural architectures for classification and domain-specific embeddings. Finally, since our experiments suggested that the decision section contains useful information, but in a realistic setting it would not be available at testing time, we would like to focus our efforts on integrating its use during the learning process.

10. Future Work

The first and foremost task of ADELE partners with regard to the continuing development of the pilot tool is to fully complete the key functionalities of automated argument extraction and outcome prediction, as well as improve annotated data and ML models as explained in Section 9.2.

Moreover, further efforts are needed also in order to improve existing and add new features based on citation and network analysis. Improvements of the present results of the models for automated extraction of key terms and case summaries, especially for Italian decisions, are also of utmost importance.

Third, a substantial part of the future activities will be aimed at completing the development of the following three functionalities, which ADELE partners had estimated to be only ready for the final version of the tool:

- *Search for similar cases* – users can ask the tool to deliver a list of decisions on factual or legal issues that are similar to the issues discussed in a decision opened by the user;

- *Search for similar arguments* – users can ask the tool to deliver a list of arguments in other decisions that are similar to a selected argument (or several consecutive arguments);

- *Visual presentation of ontology concepts* – users can navigate the ontology and its links to different parts of the decisions.

In conclusion, it can be stated that the demo version of ADELE pilot tool has fulfilled its main goal to illustrate the provisional functionalities of the online platform so that they can be demonstrated at the planned two preliminary test events with judges in Italy and Bulgaria. In addition to the remaining tasks, the opinions and recommendations to be received from the participants in these events will also be taken into account by the development of the fully functional final version of the tool.
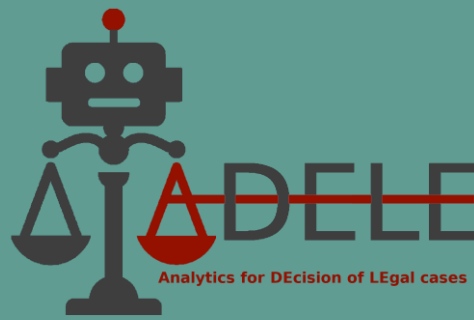
References

I. Chalkidis, M. Fergadiotis, P. Malakasiotis, N. Aletras, and I. Androutsopoulos, Legal-BERT: The muppets straight out of law school. In: *Findings of EMNLP 2020*, p. 2898, Association for Computational Linguistics.

M. Grootendorst, *KeyBERT: minimal keyword extraction with BERT*, v0.1.3. Zenodo, 2020

G. Grundler, P. Santin, A. Galassi, F. Galli, F. Godano, F. Lagioia, E. Palmieri, F. Ruggeri, G. Sartor, and P. Torroni. Detecting arguments in CJEU decisions on fiscal state aid. In: *ArgMining 2022: 9th Worhshop on Argument Mining*, 2022, accepted paper.

T. Lai, H. Ji, and C. Zhai, BERT might be overkill: A tiny but effective biomedical entity linker based on residual convolutional neural networks. In: *Findings of the Association for Computational Linguistics: EMNLP 2021*, p. 1631, Punta Cana, Dominican Republic. Association for Computational Linguistics.

N. Reimers and I. Gurevych, Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In: *EMNLP/IJCNLP* 1, 2019, p. 3980, Association for Computational Linguistics.