

Access to Justice through AI

Marco Billi and Alessandro Parenti

1 Information Technologies in Legal Practice¹

In the European Union (EU) legal system, characterized by its multilevel structure, legal professionals are required to have a deep understanding of both national and EU regulations to navigate this complex web. This intricate system is further complicated by numerous international conventions, bilateral agreements, and national instruments with overlapping rules.

In our increasingly interconnected world, legal professionals must account for the interplay of multiple normative systems, especially in international or European matters, and social interactions involving different countries. These intersecting normative systems can introduce conflicts or differing consequences, making it challenging for legal experts to determine which sources of law to consider when preparing a case. Due to the EU's multilevel legal framework, identifying the applicable legal provisions for a specific case can be extremely demanding.

In this context, Artificial Intelligence (AI) can assist users in retrieving legal information, identifying and applying relevant EU and national regulations, and assessing the degree of harmonization of national laws with the EU *acquis*. To support legal experts in these tasks, we will use a range of technologies, from expert systems designed to model legislation to large language models for user interaction.

The focus of this chapter is the *Facilex Advisory Module*, a system built on a computable version of the law, on which an expert system can be based on. By translating legal norms into rules written in a programming language, the system can apply legal reasoning to the rule base. The objective of the modeling is therefore the creation of a reasoning platform that can assist any user in finding and applying the relevant rules, by understanding which conditions have been met and which factual elements have influenced the decision.

The tool has been designed to assist legal experts and professionals in resolving practical cases involving the application of EU procedural directives

¹ The article is a joint reflection of the authors, but its drafting is broken down as follows: Billi §§ 1; 1.1; 2.1; 2.2; 2.5; 3.3; 4 Parenti §§ 1.2; 1.3; 2.3; 2.4; 3.1; 3.2; 5.

and mutual recognition instruments, as it enables legal practitioners to define a cross-border case by answering a series of questions and receiving a customized legal assessment on available legal remedies and potential cooperation issues within the specific legal system.

Below we introduce what are the main approaches taken to develop Artificial Intelligence technologies in the legal domain and provide an overview of the strengths and challenges of each one of them.

1.1 *Symbolic AI*

AI has traditionally been developed using a symbolic approach to represent the entirety of knowledge at its core. Such systems (expert systems) consist of two main parts. The first one is a knowledge base developed using methods such as logic or ontologies (i.e., symbolic methods). This represents, in a computable form, large amounts of information about the specific field the system is developed for. Most importantly, the knowledge features large sets of predefined rules that the system applies to carry out its reasoning (*rule-based systems*).²

The second part is an inference mechanism. This component represents the system's engine that enables the application of the rules contained in the knowledge base to the factual scenario provided by external users. Using both components, the system can carry out automated reasoning.

Symbolic systems experienced widespread employment in the legal domain and specifically those based on logic because of its suitability to represent legal rules.³ Legal rules, in fact, often take the form of conditional statements: "If certain conditions are met, *then* a specific legal consequence follows." This if-then structure aligns perfectly with logical implication, making it straightforward to represent legal rules using logical constructs.

An illustrative example of such symbolic systems in the legal domain is the CROSSJUSTICE (CJ) project.⁴ This project employs rule-based systems to represent and provide automated reasoning about the European laws related to the rights of suspects and accused in criminal proceedings. Both the European directives as well as several national laws applying them were transposed using logic programming. In the final platform, the system can be interrogated to know which European and national rules apply to specific scenarios provided.

2 Jürgen Karpf, 'Quality assurance of legal expert systems' (1989) *Artificial Intelligence and Law*.

3 Marek J. Sergot and others, 'The British Nationality Act as a Logic Program' (1986) 29 *Communications of the ACM* 370.

4 Crossjustice project (2022) <<https://site.unibo.it/cross-justice/en>> accessed 18/04/2024.

One of the main strengths of symbolic AI systems is the transparency and justifiability of the reasoning, thus making their overall functioning explainable.⁵ Being based on precise rules defined by humans, the process leading to the final output can be easily audited. This means that, in case of a wrong output, the process can be traced back to identify where the system made a mistake.

However, despite their strengths, rule-based systems present notable challenges, particularly in communicating their outputs to non-experts.⁶ The high-level programming languages used to express these logical conclusions often result in outputs that are difficult for laypeople to understand. This complexity can hinder the broader accessibility and usability of these systems, making it hard for non-specialists to appreciate the reasoning behind the decisions. From this perspective, the newer ML techniques and LLM can provide many advantages.

1.2 *Data Driven AI*

A different approach to the development of AI systems is represented by Machine Learning (ML). Rather than on sets of specific rules, systems are based on learning algorithms that extract patterns, i.e. learn, from large amounts of data. Several learning models were developed and implemented to obtain the best results in terms of learning outcomes.⁷ Among these, we can mention the Neural Networks (NN) model. NNs consist of large number of nodes (neurons) interconnected with each other, whose functioning is inspired to those of the neurons found in human brains. It can process the data received as input into meaningful information and propagate it to the entire network for it to learn. Deep learning, a subset of machine learning involving NNs with multiple hidden layers, has significantly advanced the capabilities of neural networks. Deep neural networks (DNNs) are particularly adept at handling high-dimensional data and extracting hierarchical representations, making them suitable for complex tasks that require multiple levels of abstraction.⁸

5 Katie Atkinson, T.J.M. Bench-Capon and Danushka Bollegala, 'Explanation in AI and law: Past, present and future.' (2020) *Artificial Intelligence* 289.

6 Harry Surden, 'Artificial intelligence and law: An overview' (2019) *Georgia State University Law Review* 35.

7 Tala Talaei Khoei, Hadjar Ould Slimane and Naima Kaabouch, 'Deep Learning: Systematic Review, Models, Challenges, and Research Directions' (2023) 35 *Neural Computing and Applications* 23103.

8 OA Montesinos López, Abelardo Montesinos López and José Crossa, 'Fundamentals of Artificial Neural Networks and Deep Learning' (2022) *Multivariate Statistical Machine Learning Methods for Genomic Prediction* 379.

In the most recent years, this approach showed a great development specifically in the field of Natural Language Processing (NLP). These techniques, when applied to the legal domain, enable the extraction and interpretation of legal concepts from vast amounts of textual data. Furthermore, they enhance the accessibility of legal information for non-experts by translating complex legal and technical jargon into more user-friendly language.

AI technologies trained to process and generate natural language, developed at a faster-ever increasing scale and have recently exploded in popularity with the creation of Large Language Models (LLMs). These powerful systems of generative AI trained on huge amounts of natural language datasets can understand and answer questions in a human-like form. Similar capabilities allow for a wide range of applications. In legal applications, LLMs have proven useful for tasks ranging from retrieving legal information⁹ to drafting contracts and legal documents, to predict judgement predictions.¹⁰ Being designed to understand and generate human language, their output is, by nature, accessible and intelligible for humans, and non-IT experts. This feature results particularly crucial in the legal domain, where connection with laypeople is a necessary requirement.

All these advantages, however, do not come without drawbacks. The other side of LLMs, just like most other data driven models, is that they are intrinsically probabilistic. This means that results can never be verified and trusted at one hundred percent, but they always entail a certain degree of mistake or, in general, non-determinism.¹¹ A direct consequence is the phenomenon of hallucinations,¹² i.e., when the model provides information that is false, non-real, or otherwise incorrect.

-
- 9 Aniket Deroy, Kripabandhu Ghosh and Saptarshi Ghosh, 'Ensemble Methods for Improving Extractive Summarization of Legal Case Judgements' (2023) *Artificial Intelligence and Law*. Andrew Blair-Stanek, Nils Holzenberger and Benjamin Van Durme, 'Can GPT-3 Perform Statutory Reasoning?' (2023) *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*.
 - 10 Dietrich Trautmann, Alina Petrova, and Frank Schilder, 'Legal prompt engineering for multilingual legal judgement prediction.' (2022). Zhongxiang Sun, 'A short survey of viewing large language models in legal aspect' (2023) <<https://arxiv.org/pdf/2303.09136>> accessed 14 July 2024.
 - 11 Bartosz Brożek, Michał Furman, Marek Jakubiec and Bartłomiej Kucharzy, 'The Black Box Problem Revisited. Real and Imaginary Challenges for Automated Legal Decision Making' (2023) *Artificial Intelligence and Law*.
 - 12 Matthew Dahl and others, 'Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models' (2024) *Journal of Legal Analysis* 64.

Moreover, in contrast to rule-based approaches, these systems are opaque by design and don't allow a verification of the process that led to the final output. The unreliability and nonexplainable nature of LLMs undermines the trust in their outputs, especially in critical legal applications where the motivations behind decisions are crucial.

1.3 *A Hybrid Approach*

Conscious of the strengths and constraints just described, we decided to take a hybrid approach and leverage the advantages of both AI paradigms.¹³

By integrating LLMs into legal expert platform we can address the communicability issues inherent in rule-based systems. While the latter excels in formalizing law with justifiable precision, LLMs enhance the usability of these systems by ensuring that their outputs are understandable to a broader audience. This symbiotic relationship between rule-based logic and LLMs creates a comprehensive tool that not only maintains the rigor of legal formalization but also makes it accessible and practical for all users, expert and layperson alike.

LLMs, with their advanced generative AI capabilities, can facilitate the incorporation of extensive domain knowledge into the system while maintaining the advantages of reliable legal reasoning and explanation inherent to logic-based methods.

The following paragraphs will first introduce the technologies applied to this project and explain how they were utilized for different tasks. In the second paragraph, we will outline the general legal principles behind the formalization of the law and their application to the current domain. The third paragraph will describe the integration of LLMs in the architecture and the conclusive paragraph will explain how everything was linked together in the available platform.

2 *The Modelling of Norms*

2.1 *General Principles of Modelling*

In this section, we will showcase the development of the rule-base layer of the Facilex advisory module, which provides the automation of reasoning within the entire system.

¹³ Marco Billi, Alessandro Parenti, Marco Sanchi and Giuseppe Pisano, 'A hybrid approach for accessible rule-based reasoning through large language models' (2024) Proceedings of JURISIN.

Among the legislation under study in the project, we concentrated on the provisions concerning the grounds for refusal of EAW, EIO, and EU freezing and confiscation orders. We chose these specific provisions due to their suitability for automation and their practical utility compared to other provisions. These rules consist of well-defined conditions for applying a norm, making them ideal for direct user queries.

We started the development of the rule-based system by building a computable representation of the targeted provisions following the methodology established for the CrossJustice project.¹⁴ For this reason, we have used the same programming language, Prolog.

During this discussion, we will focus on the relevant legal requirements to be considered while modelling different legal sources, which have first been introduced as far back as 1992.¹⁵

- *Isomorphism*, defined as one-to-one correspondence between norms in the formal model and natural language. This is achieved by representing each paragraph from each article individually and visually separated from the others, so that all grounds of refusal are kept distinct. For example, let's look at the formalization of Art 3(3) of Decision 2002/584, which states that Member States shall refuse to execute the EAW “if (1) *the person who is the subject of the European arrest warrant* (2) *may not, owing to his age, be held criminally responsible for the acts on which the arrest warrant is based under the law of the executing State.*” This specific paragraph is represented in Fig. 1 by a specific and corresponding prolog rule that is identified with the label of the article number [article3_3].

- In lines 1–3 of Fig. 1, we can see that there are two conditions for applying the article. First, that the matter concerns the EAW, second that the person is not considered criminally responsible. There are two sub rules connected, although only one is shown.

The first sub-rule (*eaw_matter*) is verified when certain conditions are fulfilled, such as declaring that the offence is covered by arts 2(2) or 2(4), that there exists an issuing authority, and an executing authority, and that the proceedings fall under the scope of Art 1. This is not shown in Fig. 1.

14 Giuseppe Contissa, Giulia Lasagni, Michele Caianiello and Giovanni Sartor (eds), *Effective protection of the rights of the accused in the EU directives: a computable approach to criminal procedure law* (Brill, 2022).

15 T.J.M. Bench-Capon and Frans P. Coenen, ‘Isomorphism and legal knowledge based systems’ (1992) 3 *Artificial Intelligence and Law*.

```

mandatory_refusal(article3_3, ExecutingMemberState,
europeanArrestWarrant):

    eaw_matter(PersonId, IssuingMemberState, ExecutingMemberState,
Offence),

    person_status(PersonId, under_age, ExecutingMemberState).
person_status(PersonId, under_age, ExecutingMemberState):-
    person_age(PersonId, Age),
    Age < 14,
    executing_member_state(ExecutingMemberState).

```

FIGURE 1 Formalization of art.3(3) Decision 2002/584

The second condition (*person_status*) is shown from lines 4–7, and states that a person is considered criminally responsible when the age is above the threshold chosen by the executing member as the age of criminal responsibility.

Therefore, although multiple sub-rules are considered, all are referred to a single paragraph of Art 3 to preserve isomorphism.

- *Reification*: rules representing legal norms need to be treated as objects that feature different properties regulated by other rules. This allows different aspects of legal systems such as jurisdiction, authority or temporal properties. In the Facilex representation, such properties are formalized only when specifically required by the norm.

Let's consider the provision of Art 4(4) of Decision 2002/584. This provides that executing authorities may refuse to apply an EAW “*where the criminal prosecution or punishment of the requested person is statute-barred according to the law of the executing Member State and the acts fall within the jurisdiction of that Member State under its own criminal law*”;

Although, generally, the jurisdiction of each state is not part of our representation, it's been explicitly stated in this article as a condition for the application of this rule. This is the only article where jurisdiction is an explicit part of the representation.

Let us now consider temporal properties of rules. We can distinguish between the external time of a rule (when a rule is in force), and its internal time (temporal issues involved in the application of a norm). The former is never explicitly formalized because there's the assumption that a rule which has been transposed is valid. The latter, instead, is captured in the modelling. For example, consider Art 18-*bis* (2-2*bis*) of the Italian Provisions to adapt national legislation to the Council Framework Decision 2002/584/JHA of 13 June 2002 on the European arrest warrant and the surrender procedures between Member States, which transposes Art 4(6) of Decision 2002/584. The article states that

```
optional_refusal(article18bis_2, MemberState, europeanArrestWarrant):-
    [...]
    person_continuous_residence(PersonId, ExecutingMemberState, Time),
    Time >= 5,
```

FIGURE 2 Partial formalization of art. 18bis(2) of Italian legislation

```
exception(optional_refusal(article4a_1_a, ExecutingMemberState,
europeanArrestWarrant), _)
```

FIGURE 3 The formalization of exceptions

When the European arrest warrant has been issued [...], the Court of Appeal may refuse to surrender of the Italian citizen or of a person who lawfully and effectively resides continuously for at least five years on the Italian territory, [...]. In this case, the period of five years must be specifically considered, and the result is the following:

In Fig. 2 you can see a partial formalization of the article. The internal time is formalized as a variable, “Time”, and must possess a value greater than 5, showing that at least 5 years of continuous residence must have passed.

- Finally, we will consider *defeasibility*. Briefly, this means providing the means to solve conflicts between rules. Rules, in fact, may have incompatible legal effects, that can be resolved through rule priorities, such as: *lex specialis* (which gives priority to the most specific rule), *lex superior* (which gives priority to the rule from the higher authority), *lex posterior* (which gives priority to the rule enacted later). The same logic applies to exclusionary rules, which are rules that provide a way to explicitly undercut other rules, making them inapplicable. We represent it by employing the Prolog based syntax for defining negation as failure. In our formalization (Fig. 3), this is achieved by having a specific predicate, called “exception” which is aimed at defining all those cases where a rule may not be applied.

For example, the conditions in Arts 4(a)(1)(a) and following in Decision 2002/584 are exceptions to the application of Art 4(a)(1). From the technical perspective, exceptions are represented as meta rules containing other rules. When an exception is valid, it prevents the contained rule from succeeding.

- The final legal requirement, which has been introduced by us due to the scope of this project, is “*harmonization*.” In the following section, we will explore in detail how European and national legal sources have been formalized and combined for a better outcome, to provide the user with an overview of the applicable grounds for refusal.

2.2 Framework Decision 2002/584

For this section, we refer to the EU provision in Fig. 1 above and showcase how that is transposed under Italian and French law.

As shown in Fig. 4 below, the transpositions feature a general rule reflecting the general EU decision provision, as well as an additional rule where the age of criminal responsibility is introduced. As we can see from the computable representation, the age of criminal responsibility is different in between France and Italy (14 in Italy, 13 in France).

It is interesting to note that, whenever the conditions allow it, 2 articles related to the same ground of refusal may simultaneously apply. Take for example a citizen of age 13, who is the subject of a European arrest warrant from France to Italy. In that case, our system will apply the Art 3(3), as well as the Italian transposition, thus showing that Italy has perfectly implemented the European article. If the issuing and executing state were reversed, as the age of criminal responsibility is respected in France, no grounds of refusal apply.

The process of harmonization, when this legal framework is considered, is grounded on the fact that the European Union leaves areas where national states are free to legislate, such as setting the age of criminal responsibility. This can be achieved by either (a) having an additional rule formalizing the conditions each state applies (as seen in the example above); or (b) by having

```
mandatory_refusal(article18_1_c, italy, europeanArrestWarrant):-
    eaw_matter(PersonId, IssuingMemberState, italy, Offence),
    person_status(PersonId, under_age, italy).
person_status(PersonId, under_age, italy):-
    person_age(PersonId, Age),
    Age < 14,
    executing_member_state(Italy).

mandatory_refusal(article695_22_3, france, europeanArrestWarrant):-
    eaw_matter(PersonId, IssuingMemberState, france, Offence),
    person_status(PersonId, under_age, france).
person_status(PersonId, under_age, france):-
    person_age(PersonId, Age),
    Age < 13,
    executing_member_state(france).
```

FIGURE 4 French and Italian transposition of art.3(3) Decision 2002/584

a predicate that the user must select to identify an event in the executing member state, when the formalization results too complex to be captured. For example, Art 4(4) of Decision 2002/584 states that an optional ground of refusal applies “*where the criminal prosecution or punishment of the requested person is statute-barred according to the law of the executing Member State and the acts fall within the jurisdiction of that Member State under its own criminal law*”.

In this case the highlighted information is formalized as a single predicate, which the user must select [`executing_proceeding_status(Offence, ExecutingMemberState, statute_barred)`]. Another way to achieve the same result would have been to have additional rules which calculate the time passed for the crime committed, and then compare it to the legal requirement for each offence, in each of the Member States. We thought that such formalization would have been too time consuming and falls outside the scope of the formalization.

2.3 Directive 2014/41

To showcase the formalization of Directive 2014/41, let us take a look at Art 11(1)(c), which states that Member States may refuse to execute the European Investigation Order when “*the EIO has been issued in proceedings referred to in Article 4(b) and (c) and the investigative measure would not be authorized under the law of the executing State in a similar domestic case*”.

The first difference we can highlight, between the Decision and the Directive, is that we have changed one of the main variables needed to apply the EIO. When verifying whether the matter concerns the application of the Directive, we ask the user to insert the “Measure” needed by the issuing Member State, instead of the offence committed by the person as we do for EAWs. As the focus of the legal source is different, we decided to reflect the same change in our formalization, as the alleged offence is accessory, and the requested measure is the main objective of the EIO.

Let us now showcase the way the connection between European sources and national sources has been formalized in our representation. As we can see from Fig. 5, the directive often leaves room for Member States to decide on their own when a measure is authorized [`national_law_does_not_authorize`].

Let us consider, for example, the case of Portugal. According to Arts 188 and 189 of the code of criminal procedure, in case of interception of telecommunications, the judge is the only authority that may authorize the measure.

This is transposed in our formalization as shown below.

It is important to note that the goal of this rule is not the ground of refusal itself, but one of the conditions for the application of the ground of refusal

```

optional_refusal(article11_1_c, ExecutingMemberState, europeanInvestigationOrder):-
    eio_matter(IssuingMemberState, ExecutingMemberState, Measure),
    (
        art4_b_applies
    ;   art4_c_applies
    ),
    national_law_does_not_authorize(ExecutingMemberState, Measure),
    \+ exception(optional_refusal(article11_1_c, ExecutingMemberState,
    europeanInvestigationOrder), _).

```

FIGURE 5 Formalization of Art 11(1)(c) Directive 2014/41

```

national_law_does_not_authorize(portugal, interception_of_telecommunications):-
    eio_matter(IssuingMemberState, portugal, Measure),
    issuing_authority(interception_of_telecommunications, IssuingAut),
    validating_authority(interception_of_telecommunications, ValidatingAut),
    ValidatingAut \= judge_or_court.

```

FIGURE 6 Conditions ex. Arts 188–189 Portuguese code of Criminal Proc.

contained in the European Directive. In particular, the last three conditions showed in Fig. 6 implement the prohibition of executing interception of telecommunications under Portuguese law when the authority validating the measure is different from a judge or a court. This Portuguese provision is integrated within the system so that when that is satisfied, it automatically triggers the application of Art 11(1)(c) of the EIO D.

Such a system enables a direct connection between national and European rules, favoring a harmonized view of the entire framework. By aligning the execution of the EIO with the specific legal conditions of each Member State, the formalization ensures that national legal standards are respected while promoting a single procedural framework across the EU. This allows each state's sovereignty and legal traditions to be preserved. Additionally, it mitigates potential conflicts by preemptively addressing discrepancies, thus showing a single outcome, instead of multiple, when the national rule is explicitly called by the European directive.

2.4 *Regulation 1805/2018*

As far as Regulation 1805/2018 is concerned, the discussion is brief. As the legal source is directly applicable in the Member States, there should theoretically be no issues in the harmonization of the Regulation.

Nonetheless, the same topics presented in the earlier chapters still apply. The regulation often sets as a condition a matter to be evaluated by the judicial authority of the executing member state, such as the principle of *ne bis in idem*. The formalization follows the same techniques presented above, by either adding rules contained in the member states, when requested by the regulation, or expanding the meaning/definition of certain matters, when national laws differ.

Something that should be clarified is that there's one condition, set by all European legal sources, which we found to be particularly troublesome. The ground for refusal applying when "*the conduct in connection [to the instrument] does not constitute a criminal offence under the law of the executing State*" is present in all three sources (EAW FD, EIO D, Regulation 1805/2018) and has been transposed as a simple "national_law_not_offence(Offence, ExecutingMemberState)". This is mainly because an automatic recognition of whether a conduct constitutes an offence is a member state falls outside the scope of this formalization, as it would require a completely different set of rules to be studied and collected.

The only exception to this consists in the offences contained in Art 2(2) EAW FD, Annex D EIO D and Art 3(1) Regulation 1805/2018. When an offence falls under one of these cases it is excluded from the verification of double criminality, and therefore, there is no need to verify whether the conduct constitutes a crime.

2.5 *Modelling of Case-Law (and Struggles)*

When dealing with the harmonization of the legal framework, it's essential to consider more than just European and national laws. Case law can play an important role in identifying how rules are to be applied. In particular, it is extremely useful when it specifies further requirements for the implementation in practice of EU provisions, rather than when it mechanically applies them.

Transposing case law into a harmonized legal framework, however, presents significant challenges, primarily due to the inherent complexity of extracting pertinent information from judicial decisions. Case law is context-specific and often contains varied interpretations contingent on each case's unique circumstances.

The extraction of universally applicable principles from case law is a difficult task due to the specificity and complexity of individual cases. We tasked legal experts with analyzing each case to generalize essential legal principles across different jurisdictions, especially when it relates to a ground of refusal, following the template below:

The case concerns the crime of [name of the offence] provided by [e.g. art x, criminal code] and is referred to the [EIO/EAW/REG 1805]. The countries involved in the cooperation procedure are [x] and [x]. In the case, the relevant ground for refusal(s) is/are [use the same wording of the EU legislation]. The main facts relevant to the case, that triggered the application (or potential application) of the ground for refusal are: [e.g. the person to be surrendered is a citizen/there is a proceeding on the same facts that has concluded with a final decision/etc ...]. Against all these factors, the court concluded that ... [summary of the conclusion of the court].

Following this methodology, we could find some rules that increased our understanding of the conditions necessary for the application of certain grounds of refusal. For example, Bulgarian sentence n. 21/2023 provided that: *“serving the sentence in the country of origin, namely Bulgaria, is related to the principle of humanity and humane treatment as far as the requested person will be placed in a prison near her permanent residence in order to maintain social contact between her and her family.”* We can transpose the information above as stating that, whenever the person requires to stay near his family, this will be considered as humane treatment of the person, therefore fulfilling one of the conditions for applying Art 4(6) of EAW FD. The formalization of this rule can be seen in Fig. 7.

Although this specific task yielded some interesting results, overall, it presented several limiting challenges. For instance, the variability of judicial reasoning and the intricate fact patterns that always feature a court case greatly complicate this process. The summaries, in fact, were always influenced by the expert's interpretation of the legal issues and often struggled to provide homogenous information. Moreover, the result often concerned the application of general national principles or procedural rules, which falls outside the purpose of our formalization. Another critical point is that case law on the

```
caselaw(article4_6, bulgaria, europeanArrestWarrant):-
    human_treatment(PersonId, MemberState).
human_treatment(PersonId, MemberState):-
    close_to_family(PersonId, MemberState).
```

FIGURE 7 Example of case law formalization

topic is quite limited, and extracting provisions that go beyond the mere application of national law has stymied our efforts in this task.

In the end, such manual summarization has proven to be time-consuming and potentially exacerbating the challenge of creating a harmonized legal framework, as this process demands significant human resources and expertise, limiting the efficient integration of case law into a cohesive legal system.

3 LLM Integration

3.1 *Architecture*

As outlined above, we implemented a hybrid architecture that leverages the strengths of both symbolic and sub-symbolic artificial intelligence to obtain a result that is at the same time verifiable and accessible. In particular, we employed Large Language Models to process the output produced by the rule-based system and gave it a form that is understandable to laypeople.

We use the GPT-4 model developed by Open-AI, in that it has proved to be the most effective and provides the highest percentage of positive performance.¹⁶ We used an API to enable the setting of various parameters. We set the system's temperature to the lowest possible, i.e., 0. This minimizes the system's tendency to be creative and generate words or sentences that do not strictly derive from the knowledge provided by the user.

At the center of the methodology is the creation of a prompt carrying all the necessary information and instruction needed by the LLM to carry out its task. This includes two main elements.

First, the output of the rule-based system. This consists of a prolog tree indicating the rules that were applied to solve the specific scenario. As described in section 2, the Facilex automated reasoning system focuses on the regulation of the grounds of refusal of EAW, EIO, Freezing Order (FO) and Confiscation Order (CO), as defined by EU legislation. Therefore, the output of the system describes if and which ground of refusal applies together with all the satisfied conditions that led to the application. This is presented in a prolog tree indicating the rules applied, written in the prolog programming language.

The second part of the prompt features the command given to the LLM. The task is described, giving exact instructions on the content and form requested, and on the procedure to be followed to carry it out. As anticipated, our goal is

¹⁶ See <<https://huggingface.co/spaces/lmsys/chatbot-arena-leaderboard>> accessed 18 June 2024.

to obtain a transposition (or explanation) in natural language of the of the outcome provided by the rule-based system, thus allowing understanding for end users regardless of their IT expertise. This means providing the main result, i.e., whether and which ground of refusal apply. Moreover, in order for this explanation to be exhaustive and effective in giving a full understanding of the legal situation to the users, it should also include the reasoning behind the application of the ground of refusal in connection with the specific scenario provided.

3.2 *Prompt Engineering and Challenges*

We encountered several issues. An initial challenge was to obtain a fixed structure in the output, which was often presented in different formatting (e.g., bullet lists, numbered lists, free-form text). We found that it was necessary to give a fixed structure for the LLM to follow, in order to decrease the degree of variance between answers and maximise repeatability and reliability.

Second, we noticed a propensity to miss and not mention key facts part of the specific scenario. To address this, we included in the prompt the requirement for the LLM to take into consideration all legal terms and facts from the scenario, and to refer to those when building the explanation.

Finally, the most challenging aspect was to overcome the tendency of the model to ascribe meaning to legal text where not explicitly mentioned nor provided. Often, meaning was guessed based on the most general significance, disregarding the specificity of juridical norms and legal lexicon.

Trying to address all these challenges we engineered the prompt shown in Fig. 8. The LLM is provided with a clear structure to follow (Summary, What, Why). Moreover, the prompt stresses the importance of referencing all the Prolog terms found in the tree, which also include the part of the factual scenario.

```
You have been provided a Prolog inference tree using a legal norm in a specific
case (Prolog Tree). Provide the following info according to the given structure:

- Summary : simplified text of the legal norm . Use everyday language with a
serious register . Add any information you have on the source of law;
- What grounds of refusal apply : what kind of grounds of refusal apply
(mandatory or optional), if any, according to the Prolog explanation;
- Why do they apply : inference steps and reasoning that led to the grounds of
refusal. Use all the Prolog terms in the explanation explicitly referencing
the original Prolog when needed. Ignore empty Prolog variables (_number).

Use enumerations in the ' What grounds of refusal apply ' and ' Why do they apply '
sections if needed.
```

FIGURE 8 Natural language explanation prompt

3.3 *Implementation*

For exemplification purposes, we will now demonstrate the functioning of the LLM integration using a real-world scenario. Consider the following legal case inspired by a real court case: Seydou is a boy born in France on January 1, 2010, to Ivorian parents and residing in Bobigny. On the evening of December 31, 2023, together with a group of friends, he committed the crime of theft of spirits. After being caught and temporarily released by the French police under the custody of his parents, Seydou was taken to Italy by his parents to distance him from bad company. The French authorities, having identified Seydou through security cameras, issued a European arrest warrant.

Such a case can be submitted to the Facilex platform to identify which, if any, grounds of refusal may apply to the EAW issued for Seydou. After issuing all the necessary information related to the case, the system will return an outcome as shown in Fig. 9. This indicates the grounds for mandatory refusal according to Art 3(3) of EAW FD. It also includes a Prolog tree that outlines all the inference steps taken by the system to reach this conclusion.

Clearly, however, the output is not intelligible to the average user. While it mainly features natural language terms, these follow a form and syntax specific to the Prolog programming language, which is opaque to non-IT experts.

To improve the readability of this output and make it accessible to laypeople, we feed it to an LLM along with the prompt described above.

The resulting output is notable and shown in Fig. 10. The LLM organizes its output according to the given structure: first, it provides a summary of the applicable norm; second, it indicates the type of refusal ground; third, it details all the inference steps that led to the application of this refusal ground. In doing so, the LLM explicitly references the Prolog terms present in the original output, as requested, enabling the verification of the answer. This last part is particularly significant as it allows the user not only to understand the applied provision but also how the specific elements of the scenario influenced the outcome. Lastly, the answer included a summary of the legal case and its solution at the end.

In this specific case, Art 3(3) of EAW FD finds application. This provides that a European arrest warrant must be refused when the person who committed the crime cannot be considered criminally responsible under the law of the executing state. In our scenario, Seydou was 13 when he committed the crime, while the threshold for criminal responsibility according to Italian law is 14.

To conclude, the LLM respected the guidelines given as to the form of the output, regarding both the structure and the lexicon of the answer. Moreover, the answer is also juridically correct. On the one hand, it reports all the information present in the Prolog output and, most importantly, it does not include

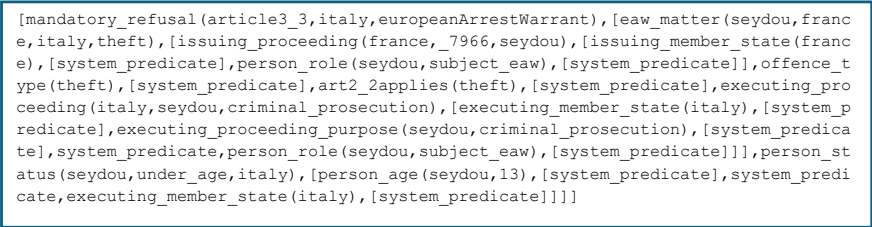


FIGURE 9 Facilex platform's output

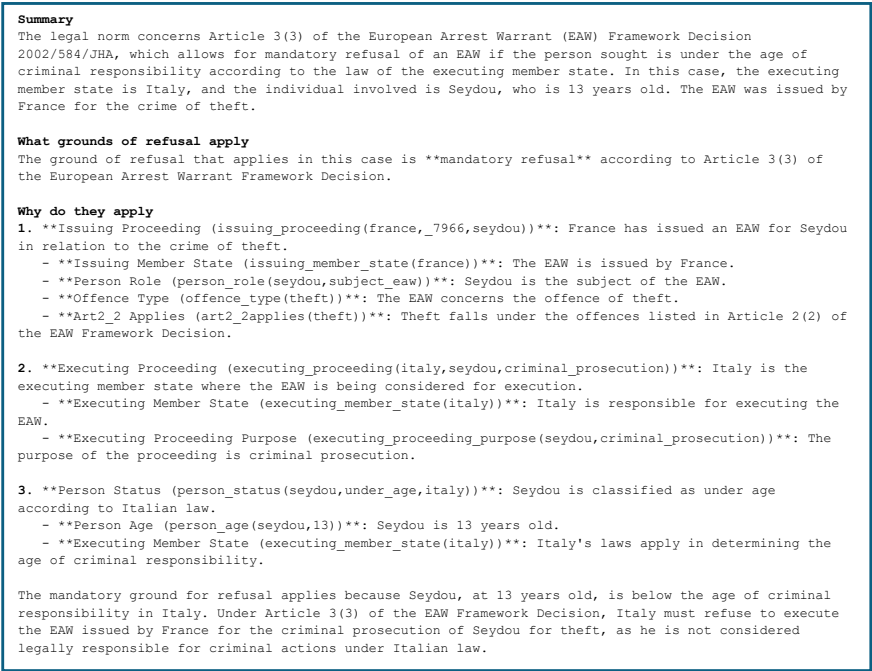


FIGURE 10 LLM's transposition to natural language

any self-generated facts or information. This methodology proved successful in improving accessibility to the Facilex system, while preserving the soundness and trustworthiness of the underlying rule-based reasoning.

4 Platform Integration

4.1 User Interface

To conclude, we will demonstrate the practical application of this system. As previously discussed, we have formalized several rules in a programming

language to represent European and national legal sources. However, executing these rules successfully requires specific conditions to be met, which is where the user interface becomes essential. By accessing the Facilex Advisory Module through this link¹⁷ users can input all relevant facts related to their scenario by answering the platform’s questions. Based on this information, the tool can analyze all relevant European and national legal norms and present inferred solutions with a detailed assessment. Each question is linked to a single Prolog fact, and, through the answer, the fact is adjusted to be tailored to the user.

The questions have been (mostly) taken from the annexes provided by the European Union, at the foot of the legislative document. These represent the official form that authorities must fill in when issuing the request for the instrument. To these questions, we have added others which are essential for applying the specific grounds of refusal.

The first questions are general, and are used to determine which further questions should be displayed to the user:

FA

CI

LEN

FACILEX ADVISORY MODULE

CUSTOMISED ASSESSMENT

What are you interested in?

Select

▼

What is the issuing State?

Select

▼

What is the executing State?

Select

▼

Submit

Clear

FIGURE 11 User interface: general questions

The user shall select which legal instrument he’s interested in, and then select the issuing and executing state. As these questions should be answered every time the platform is called upon, they are mandatory. All other questions are optional, although failure to provide answers may result in the platform being unable to generate a positive outcome.

¹⁷ See Facilex Advisory Module <<https://facilex-tool.eu/Advisory/customised-assessment>> (2024) accessed 18 June 2024.

If we were to select “European Arrest Warrant” as the instrument, the following questions would be asked, which represent points (a) and (e) of the Annex on EAW.

In some cases, the information is inserted as free text (e.g, the person’s name). Other times, instead, the user is restricted to a multiple-choice answer, such as in the representation of point (e) of the Annex. This is because, while the person’s name can have a random value, the list of offences must be confined to those analyzed by legal experts.

To elaborate, a variable can be substituted with any value and remain valid. However, there are rules that apply only to specific offences and are included in our formalization’s logic. This makes the precise naming of the offense crucial for our purpose.

Information regarding the identity of the requested person: Name and Forename

Information regarding the identity of the requested person: Nationality

Select

FIGURE 12 User interface: annex point (a)

Offences: Nature and legal classification of offence, and the applicable statutory provision

Select

Offences: State in which the offence was committed

Select

FIGURE 13 User interface: annex point (b)

(a) Information regarding the identity of the requested person:

Name:

Forename(s):

Maiden name, where applicable:

Aliases, where applicable:

Sex:

Nationality:

FIGURE 14 Annex point (a)

(e) Offences:

This warrant relates to in total: offences.

Description of the circumstances in which the offence(s) was (were) committed, including the time, place and degree of participation in the offence(s) by the requested person:

.....

.....

.....

Nature and legal classification of the offence(s) and the applicable statutory provision/code:

.....

.....

.....

.....

FIGURE 15 Annex point (e)

Is the offence covered by amnesty in the executing Member State?

☐ Yes

☐ No

FIGURE 16 User interface: amnesty question

Additionally, we include questions beyond those provided in the annex, which are necessary for the executing member state to determine whether any grounds for refusal exist. For example, the user may be asked to consider if the offense in question is covered by amnesty.

The intended user of this platform is the authority in the executing member state who receives the EAW (or another instrument) and is deliberating on its execution. There are certain grounds that apply only based on questions connected to the annex, while others need additional information to be provided, such as amnesty.

One limitation of this approach is that the system heavily relies on the legal knowledge of the user to answer certain questions. While a few conditions are straightforward, most require at least some familiarity with legal terminology and concepts. This restricts the user base to those with relevant expertise, such as authorities involved in drafting the instrument or lawyers seeking support for their clients.

Finally, once the system evaluates the outcome, the solution presented in §3.3 is displayed.

5 Concluding Remarks

In this chapter we have shown the results obtained by combining two approaches in the domain of legal informatics. The Facilex project integrates

these approaches to address the formalization of grounds of refusal in the three mutual recognition instruments, effectively applying legal rules and explaining them in natural language. This project harnesses old and new technologies to perform various tasks that are essential to legal informatics.

Firstly, we introduced the technologies applied in this project, explaining their use for different tasks. Subsequently, we outlined the general legal principles behind the formalization of the law and demonstrated their application in the current domain. Furthermore, we described the integration of LLMs within the architecture, highlighting their role in enhancing the project's capabilities.

The Facilex project exemplifies reliability and trustworthiness, standing out as an innovation in the field. This integrated approach not only ensures accuracy and precision but also enhances accessibility and comprehension for users. Additionally, this effort is essential in harmonizing the cooperation of states in mutual recognition instruments, facilitating smoother cross-border legal interactions. In conclusion, the Facilex platform represents a significant advancement in legal informatics, offering a robust and innovative solution for the formalization and application of legal principles.