

Beyond Logic Programming for Legal Reasoning^{*}

Ha Thanh Nguyen^{1,*}, Francesca Toni², Kostas Stathis³ and Ken Satoh¹

¹National Institute of Informatics (NII), 2-1-2 Hitotsubashi, Chiyoda City, Tokyo, Japan

²Imperial College London, Exhibition Rd, South Kensington, London SW7 2BX, United Kingdom

³Royal Holloway University of London, Egham Hill, Egham TW20 0EX, United Kingdom

Abstract

Logic programming has long been advocated for legal reasoning, and several approaches have been put forward relying upon explicit representation of the law in logic programming terms. In this position paper we focus on the PROLEG logic-programming-based framework for formalizing and reasoning with Japanese presupposed ultimate fact theory. Specifically, we examine challenges and opportunities in leveraging deep learning techniques for improving legal reasoning using PROLEG, identifying four distinct options ranging from enhancing fact extraction using deep learning to end-to-end solutions for reasoning with textual legal descriptions. We assess advantages and limitations of each option, considering their technical feasibility, interpretability, and alignment with the needs of legal practitioners and decision-makers. We believe that our analysis can serve as a guideline for developers aiming to build effective decision-support systems for the legal domain, while fostering a deeper understanding of challenges and potential advancements by neuro-symbolic approaches in legal applications.

Keywords

PROLEG, Deep Learning, ProbLog, Prolog, Neurosymbolic, Legal Reasoning

1. Introduction

Legal reasoning is a complex process that involves evaluating and applying principles, rules, and regulations from various sources, such as legislation, case law, and general legal principles. The advancement of AI, particularly in terms of deep learning and natural language understanding, has created new opportunities for developing systems that can aid legal practitioners in navigating this complex landscape.

In this paper, we examine the challenges and opportunities in leveraging deep learning techniques for improving legal reasoning using the PROLEG system [1], a logic-programming-based framework for formalizing and reasoning with the Japanese presupposed Ultimate Fact theory (*JUF theory*, in short, from *Youken-jijisturon*, in Japanese). The JUF theory is used for decision-making by judges under incomplete information.

PROLEG accommodates a representation of the JUF theory in logic programming terms, reflecting lawyers' reasoning by drawing upon the idea of "openness" proposed for the JUF

London'23: Workshop on Logic Programming and Legal Reasoning, July 09–10, 2023, London, UK

*Corresponding author.

✉ nguyenhathanh@nii.ac.jp (H. T. Nguyen); f.toni@imperial.ac.uk (F. Toni); Kostas.Stathis@rhul.ac.uk (K. Stathis); ksatoh@nii.ac.jp (K. Satoh)

🆔 0000-0003-2794-7010 (H. T. Nguyen); 0000-0001-8194-1459 (F. Toni); 0000-0002-9946-4037 (K. Stathis); 0000-0002-9309-4602 (K. Satoh)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

theory. However, the deployment of PROLEG to reason about particular input cases requires the addition of suitable facts, relating to the cases, to the system.

In this position paper we explore various options for integrating deep learning techniques into the legal reasoning process envisaged by PROLEG. We discuss these options in the context of extracting and leveraging facts from natural language texts, ranging from enhancing fact extraction to end-to-end solutions for reasoning with textual legal descriptions. We provide toy illustrations of the various options.

Our analysis is targeted at developers working on decision-support systems for legal practitioners, and our proposed options are designed to address the challenges and limitations observed in related works, such as understanding and applying prescriptive rules specified in natural language, providing better structural insights, and effectively representing and sharing legal knowledge using well-defined logical languages.

Throughout the paper, we assess the advantages and limitations of each option, considering their technical feasibility, interpretability and alignment with the needs of legal practitioners such as judges, jurors, and lawyers involved in decision-making based on legal information. By combining deep learning with symbolic reasoning methods, we aim to provide developers with insights and solutions for building effective and comprehensible decision-support systems in the legal domain.

2. Related Work

Natural Language Processing (NLP) and Natural Language Understanding (NLU) tasks applied to law have proliferated in the recent years, with several works in the literature focusing on the challenge of legal reasoning and decision-making using structured, logical representations derived from texts. This work relates to several areas of active research.

Holzenberger et al. [2] presents a dataset and legal-domain text corpus to investigate the performance of NLU approaches on statutory reasoning. The study compares the results of machine reading models with a hand-constructed Prolog-based system, highlighting the challenges facing statutory reasoning moving forward. This work demonstrates the importance of understanding and applying prescriptive rules specified in natural language for legal reasoning. Further work by Holzenberger and Van Durme [3] decomposes statutory reasoning into four types of NLU challenge problems by introducing concepts and structures found in Prolog programs. Their results show that models for statutory reasoning benefit from the additional structure, leading to improved performance over prior baselines and finer-grained model diagnostics.

Palmirani et al. [4] extends the RuleML language [5] to create LegalRuleML, enabling the effective sharing and exchange of legal knowledge between documents, business rules, and software applications. This work focuses on detecting, modeling, and expressively representing legal knowledge to support legal reasoning and its application in the business rule domain. In a similar context, Schneider et al. [6] describe the EU-funded project Lynx, which aims to create a Legal Knowledge Graph (LKG) for the semantic processing, analysis, and enrichment of legal documents. The article discusses the use cases, platforms, and semantic analysis services that operate on the documents for more effective legal information management.

Early work by Nakamura et al. [7] proposes a framework for translating legal sentences into

logical form, and describes the implementation of such an experimental system. In a similar vein, Lagos et al. [8] present an event extraction mechanism based on NLP techniques to extract the use of entity related information corresponding to the relations among the parties of a case in the form of events. More recently, the work by Gaur et al. [9] address the translation of legal text into formal representations using the NL2KR system. This system translates legal text into various formal representations, enabling the use of existing logical reasoning approaches on legal text in English. By allowing reasoning with text, this work bridges the gap between natural language processing and existing logical reasoning frameworks designed for legal information.

In the remainder, we will explore several options for improving the performance and applicability of the PROLEG legal reasoning support system [1]. Each option proposed is based on the related work and methods discussed above, aiming to enhance the reasoning capabilities of PROLEG by incorporating deep learning techniques for extracting and leveraging facts from natural language texts. These options attempt to address different challenges and limitations observed in previous works, such as the difficulty in understanding and applying prescriptive rules specified in natural language [2], the need for additional structural insights provided by Prolog-based systems [3], and the importance of effectively representing and sharing legal knowledge using well-defined logical languages [4].

3. Options

We focus here on the following setting: laws have been mapped onto PROLEG; there are text descriptions of the “facts” to which the laws need to be applied; we want to “emulate” reasoning with the laws and facts by using PROLEG and a logical description of the “facts”.¹

3.1. Option 1: PROLEG + Fact Extraction by Deep Learning

This is the approach envisaged with PROLEG, whereby the facts are standard PROLOG facts and are to be used in combination with PROLEG using its standard inference mechanism. The fact extraction is carried out by fine-tuning a language model.

Fact extraction requires carrying out several sub-tasks, given an input text (the query) describing the facts in natural language:

- Identifying which (article of) law matches the query;
- Identifying which bits of the (article of) law matches the query;
- Populate the logical facts underpinning the law with the information from the query matching the law (see [10, 11]).

¹This setting is in line with the goals of a *judge* or *juror* who needs to decide on the outcome of a case. Of course, there are other settings of interest. For example, we could be wanting to generate the “facts” so that the outcome of the application of the law is “driven” by the lawyer presenting the “facts”. In this alternative setting, we may want to look for most “beneficial facts”. This latter setting is in line with the operation of defence or plaintiff lawyers.

3.2. Option 2: ProbLEG: PROLEG + Probabilistic Fact Extraction by Deep Learning

This is an alternative approach, whereby the facts are probabilistic as in ProbLog [12] and are to be used in combination with PROLEG using ProbLog’s inference mechanism. The fact extraction needs to be again carried out externally to PROLEG and ProbLog, e.g. by fine-tuning a language model [13, 14, 15, 16].

3.3. Option 3: Deep ProbLEG

The first two approaches are neuro-symbolic in a loose sense (in that the deep learning and reasoning components interface but are kept separate). A fundamentally different approach integrates reasoning and fact extraction tightly in the spirit of Deep ProbLog [17]. This approach would require injecting PROLEG into a deep architecture of the kind explored in [17], whereby reasoning a-la-ProbLog follows probabilistic fact extraction within a cohesive whole.

3.4. Option 4: End-to-End Deep Learning

This final approach is aimed at serving as a baseline and is expected to be less interpretable than the other three, based on fine-tuning language models, such as domain-oriented models [18, 19] or robust domain-independent models [20, 21, 22], to reason directly with the textual representation of the legislation with the textual representation of the facts. This approach would not rely on reasoning with PROLEG directly and will be non-interpretable.

Methodologically, we could see the legal reasoning that PROLEG is carrying out as a form of *soft theorem proving* [23], possibly fine-tuning the adversarial method in [24].

4. A toy illustration

Here we focus on illustrating the first two options with a simple toy example with standard PROLEG structure but made-up laws. Our aim is to show what the two options may amount to. A first law may be represented in PROLEG by:

```
right(B, S) <=  
  cond(B, S).
```

```
cond(B, S) <=  
  sub_cond(B, S).
```

sub_cond may amount to ‘being friends’.

```
exception(sub_cond(B, S), ex(B, S)).
```

```
ex(B, S) <=  
  ex_cond(B, S).
```

sub_cond may amount to ‘knowing each other for a long time’.

A second law may be represented simply as:

```
other_right(B, S) <= smtg(B, S).
```

ex_cond may amount to ‘working together’

Facts are about *sub_con*(*B, S*) and *ex_con*(*B, S*), as well as *smtg*(*B, S*).

Suppose the textual description of three cases of interest is as follows:

case 1 Thanh and Kostas met recently but are good friends already (*thus they satisfy the sub-conditions for the first law*);

case 2 Ken and Francesca have collaborated for a while (*do they satisfy the conditions of the second law? the exception to the first law?*) and are in good terms (*do they satisfy the conditions of the first law?*);

case 3 Thanh is working under Ken’s supervision (*do they trigger the second law?*) and they are in good terms (*do they satisfy the sub-conditions of the first law?*).

Here, case 1 most definitely matches the first law, and case 3 probably matches the second law (but may also be deemed somewhat to match the first law); it is very unclear though whether case 2 matches the first law (with exception) or the second law. Also, after deciding which law case 2 matches, there is uncertainty in extracting the facts.

Mapping PROLEG into PROLOG. We focus here on the object-level translation of the above PROLEG program [1], as follows:

```
right(B, S) :- cond(B, S).  
cond(B, S) :- sub_con(B, S), not ex_con(B, S).  
other_right(B, S) :- smtg(B, S).
```

```
sub_con(thanh, kostas).  
sub_con(ken, fran).
```

```
ex_con(ken, fran).  
smtg(thanh, ken).
```

This translation makes use of transforming exception clauses in PROLEG as negation by failure in the condition of the rules². An alternative translation that we may want to consider is given in Appendix A.

²This is compatible with the spirit of the standard translation methods for PROLEG already described in <http://research.nii.ac.jp/ksatoh/juris-informatics-papers/jurix2009-ksatoh.pdf>

Option 1 Deep learning learns “crisp” logical facts, e.g. they may be:

`sub_con(thanh, kostas). %(first case)`

`sub_con(ken, fran). %(and)`
`ex_con(ken, fran). %(second case)`

`smtg(thanh, ken). %(third case)`

Option 2 Deep learning learns probabilities. What are these probabilities on? The facts or the applicability of the laws? This makes a difference if $(cond(B, S) \cup ex_con(B, S)) \cap smtg(B, S) \neq \emptyset$. Also, even if $(cond(B, S) \cup ex_con(B, S)) \cap smtg(B, S) = \emptyset$, the probability of matching a law may play a role in determining the probability of the facts holding.

Determining the probability of facts in isolation. This may result in something like

`0.99::sub_con(thanh, kostas).`

`0.51::sub_con(ken, fran).`

`0.73::ex_con(ken, fran).`

`0.49::smtg(ken, fran).`

`0.51::sub_con(thanh, ken).`

`0.65::smtg(thanh, ken).`

Determining the probability of facts while taking into account the relevance of laws. This means that the probability of a fact may be a function of two probabilities: that the case matches a law and that phrases in the case description fit atomic templates. For example, we may compute the following probabilities that the cases match the laws:³

- Case 1 matches law 1 with probability 1;
- Case 2 matches law 1 with probability 0.7 and law 2 with probability 0.3;
- Case 3 matches law 2 with probability 0.6 and law 1 with probability 0.4.

Then, for each case (when the law applicability is probable according to some threshold), we determine the probability of the atomic templates $sub_cond(\cdot, \cdot)$, $ex_cond(\cdot, \cdot)$, and $smtg(\cdot, \cdot)$ being instantiatable on the textual description of the cases.

³We may need to assume that the two laws are independent.

Suppose that case 1 matches $sub_cond(\cdot, \cdot)$ with probability 0.99. Then, the resulting probabilistic fact is

```
(1 * 0.99) = 0.99::sub_con(thanh, kostas).
```

Suppose that case 3 matches $smtg(\cdot, \cdot)$ with probability 0.9. Then, the resulting probabilistic fact is

```
(0.6 * 0.9) = 0.54::sub_con(ken, fran).
```

Keeping the relevance of the laws and the matching of the atomic templates separate. This amounts to treating the applicability of laws “abductively”, and may lead (in our example) to:

```
right(B, S) :- cond(B, S), applicable_1(B, S).
```

```
cond(B, S) :- sub_con(B, S), not ex_con(B, S).
```

```
other_right(B, S) :- smtg(B, S), applicable_2(B, S).
```

Then the fact:

```
0.95::sub_con(thanh, kostas).
```

```
0.55::sub_con(ken, fran).
```

```
0.78::ex_con(ken, fran).
```

```
0.42::smtg(ken, fran).
```

```
0.50::sub_con(thanh, ken).
```

```
0.69::smtg(thanh, ken).
```

```
0.99::applicable_1(thanh, kostas).
```

```
0.51::applicable_1(ken, fran).
```

```
0.49::applicable_2(ken, fran).
```

```
0.71::applicable_1(thanh, kenn).
```

```
0.81::applicable_2(thanh, kenn).
```

5. Discussions and Conclusion

We have explored challenges and opportunities of combining deep learning techniques and the PROLEG system to improve legal reasoning. Our analysis has identified four distinct

options ranging from enhancing fact extraction using deep learning to end-to-end solutions for reasoning with textual legal descriptions. Throughout our analysis, we have taken the standpoint of developers rather than legal practitioners. In other words, the systems we envision could be used as the backend of decision-support systems, potentially by enforcing certain threshold conditions to facilitate decision-making and reasoning in legal contexts.

Although we primarily focused on a "judge/juror" setting for developers, some of the analyses and methodologies we propose could be adapted to a "lawyer" setting, as discussed earlier. In this context, our systems could be used to support the selection and presentation of facts that lead to a desired conclusion (or a high probability of it) and aid in evidence retrieval to build strong legal arguments and cases.

By integrating deep learning with symbolic reasoning methods, we can develop more robust and flexible legal reasoning systems that can handle the challenges posed by natural language ambiguity and rigid logical representations.

Some additional benefits of the proposed approaches include:

- A better understanding of the strengths and weaknesses of different arguments and the underlying evidence.
- Effective identification of crucial evidence and the areas where such evidence can be found to overturn a legal decision if necessary.
- Enhanced capability to debug legal decisions and regulations, allowing for more efficient evaluation and adjustments in legal frameworks.

In conclusion, by leveraging the advancements in deep learning and natural language understanding, the proposed options aim to provide developers with practical tools and methods for building decision-support systems that aid various legal practitioners, including judges, jurors, and lawyers, in navigating the complex landscape of legal reasoning and decision-making. Moving forward, we plan to delve deeper into the four options, providing a comprehensive examination and evaluation of each method to further enhance our understanding of their potential impact and applications in the legal domain.

Acknowledgments

This work was supported by JSPS KAKENHI Grant Number JP22H00543 and JST, AIP Trilateral AI Research Grant Number JPMJCR20G4. Francesca Toni and Kostas Stathis would like to thank the National Institute of Informatics, Tokyo, Japan, for supporting their visit to Japan that made this work possible. Francesca Toni also acknowledges support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No.101020934, ADIX), as well as support from J.P. Morgan and the Royal Academy of Engineering, UK, under the Research Chairs and Senior Research Fellowships scheme.

References

- [1] K. Satoh, K. Asai, T. Kogawa, M. Kubota, M. Nakamura, Y. Nishigai, K. Shirakawa, C. Takano, Proleg: an implementation of the presupposed ultimate fact theory of japanese civil code by prolog technology, in: JSAI international symposium on artificial intelligence, Springer, 2010, pp. 153–164.
- [2] N. Holzenberger, A. Blair-Stanek, B. Van Durme, A dataset for statutory reasoning in tax law entailment and question answering, arXiv preprint arXiv:2005.05257 (2020).
- [3] N. Holzenberger, B. Van Durme, Factoring statutory reasoning as language understanding challenges, arXiv preprint arXiv:2105.07903 (2021).
- [4] M. Palmirani, G. Governatori, A. Rotolo, S. Tabet, H. Boley, A. Paschke, Legalruleml: Xml-based rules and norms., RuleML America 7018 (2011) 298–312.
- [5] H. Boley, A. Paschke, O. Shafiq, Ruleml 1.0: the overarching specification of web rules, in: Semantic Web Rules: International Symposium, RuleML 2010, Washington, DC, USA, October 21-23, 2010. Proceedings 4, Springer, 2010, pp. 162–178.
- [6] J. M. Schneider, G. Rehm, E. Montiel-Ponsoda, V. Rodríguez-Doncel, P. Martín-Chozas, M. Navas-Loro, M. Kaltenböck, A. Revenko, S. Karampatakis, C. Sageder, et al., Lynx: A knowledge-based ai service platform for content processing, enrichment and analysis for the legal domain, Information Systems 106 (2022) 101966.
- [7] M. Nakamura, S. Nobuoka, A. Shimazu, Towards translation of legal sentences into logical forms, in: K. Satoh, A. Inokuchi, K. Nagao, T. Kawamura (Eds.), New Frontiers in Artificial Intelligence, Springer Berlin Heidelberg, Berlin, Heidelberg, 2008, pp. 349–362.
- [8] N. Lagos, F. Segond, S. Castellani, J. O’Neill, Event extraction for legal case building and reasoning, in: Z. Shi, S. Vadera, A. Aamodt, D. Leake (Eds.), Intelligent Information Processing V, Springer Berlin Heidelberg, Berlin, Heidelberg, 2010, pp. 92–101.
- [9] S. Gaur, N. H. Vo, K. Kashihara, C. Baral, Translating simple legal text to formal representations, in: New Frontiers in Artificial Intelligence: JSAI-isAI 2014 Workshops, LENLS, JURISIN, and GABA, Kanagawa, Japan, October 27–28, 2014, Revised Selected Papers, Springer, 2015, pp. 259–273.
- [10] H.-T. Nguyen, F. Wachara, F. Nishino, K. Satoh, A multi-step approach in translating natural language into logical formula, in: Legal Knowledge and Information Systems, IOS Press, 2022, pp. 103–112.
- [11] M. Z. May, N. Ha-Thanh, S. Ken, S. Saku, N. Fumihito, Improving translation of case descriptions into logical fact formulas using legalcasener, in: Nineteenth International Conference on Artificial Intelligence and Law (ICAIL 2023), 2023.
- [12] L. D. Raedt, A. Kimmig, Probabilistic (logic) programming concepts, Mach. Learn. 100 (2015) 5–47. URL: <https://doi.org/10.1007/s10994-015-5494-z>. doi:10.1007/s10994-015-5494-z.
- [13] Y. T.-H. Vuong, Q. M. Bui, H.-T. Nguyen, T.-T.-T. Nguyen, V. Tran, X.-H. Phan, K. Satoh, L.-M. Nguyen, Sm-bert-cr: a deep learning approach for case law retrieval with supporting model, Artificial Intelligence and Law (2022) 1–28.
- [14] H.-T. Nguyen, M.-K. Phi, X.-B. Ngo, V. Tran, L.-M. Nguyen, M.-P. Tu, Attentive deep neural networks for legal document retrieval, Artificial Intelligence and Law (2022) 1–30.
- [15] H.-T. Nguyen, M.-P. Nguyen, T.-H.-Y. Vuong, M.-Q. Bui, M.-C. Nguyen, T.-B. Dang, V. Tran,

- L.-M. Nguyen, K. Satoh, Transformer-based approaches for legal text processing: Jnlp team-coliee 2021, *The Review of Socionetwork Strategies* 16 (2022) 135–155.
- [16] Y. Shao, J. Mao, Y. Liu, W. Ma, K. Satoh, M. Zhang, S. Ma, Bert-pli: Modeling paragraph-level interactions for legal case retrieval., in: *IJCAI, 2020*, pp. 3501–3507.
- [17] R. Manhaeve, S. Dumancic, A. Kimmig, T. Demeester, L. D. Raedt, Neural probabilistic logic programming in DeepProbLog, *Artif. Intell.* 298 (2021) 103504. doi:10.1016/j.artint.2021.103504.
- [18] H.-T. Nguyen, H.-Y. T. Vuong, P. M. Nguyen, B. T. Dang, Q. M. Bui, S. T. Vu, C. M. Nguyen, V. Tran, K. Satoh, M. L. Nguyen, Jnlp team: Deep learning for legal processing in coliee 2020, arXiv preprint arXiv:2011.08071 (2020).
- [19] I. Chalkidis, M. Fergadiotis, P. Malakasiotis, N. Aletras, I. Androutsopoulos, Legal-bert: The muppets straight out of law school, arXiv preprint arXiv:2010.02559 (2020).
- [20] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P. J. Liu, Exploring the limits of transfer learning with a unified text-to-text transformer, *The Journal of Machine Learning Research* 21 (2020) 5485–5551.
- [21] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, *Advances in neural information processing systems* 33 (2020) 1877–1901.
- [22] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann, et al., Palm: Scaling language modeling with pathways, arXiv preprint arXiv:2204.02311 (2022).
- [23] P. Clark, O. Tafjord, K. Richardson, Transformers as soft reasoners over language, in: *IJCAI, 2020*.
- [24] A. Gaskell, Y. Miao, L. Specia, F. Toni, Logically consistent adversarial attacks for soft theorem provers, *CoRR abs/2205.00047* (2022). To appear (IJCAI 2022).

A. Alternative PROLEG translations

A somewhat more sophisticated object-level translation of the PROLEG program in Section 4 could be as follows:

```

right(B, S) :- cond(B, S).

cond(B, S) :- sub_con(B, S), not exempt_sub_con(B, S).

exempt_sub_con(B, S) :-
    ex_con(B, S), not exempt_ex_con(B, S).

other_right(B, S) :- smtg(B, S).

sub_con(thanh, kostas).

```

```
sub_con(ken, fran).
```

```
ex_con(ken, fran).
```

```
smtg(thanh, ken).
```

The latter translation method is more methodological and has hooks for exceptions of exceptions.