Explainable Reasoning with Legal Big Data: A Layered Framework

Grigoris Antoniou
University of Huddersfield, Queensgate, Huddersfield, HD1 3DH, UK
g.antoniou@hud.ac.uk

Katie Atkinson
University of Liverpool, UK
K.M.Atkinson@liverpool.ac.uk

George Baryannis
University of Huddersfield, Queensgate, Huddersfield, HD1 3DH, UK
g.bargiannis@hud.ac.uk

Sotiris Batsakis
Technical University of Crete, Chania, 73100, Greece and University of Huddersfield, Queensgate, Huddersfield, HD1 3DH, UK
s.batsakis@hud.ac.uk

Luigi Di Caro
University of Turin, Turin, Italy
dicaro@di.unito.it

Guido Governatori
Independent researcher
guido@governatori.net

Livio Robaldo
Legal Innovation Lab Wales, Hillary Rodham Clinton School of Law, Swansea University, Singleton Park, Swansea, SA2 8PP, UK
livio.robaldo@swansea.ac.uk

Giovanni Siragusa
University of Turin, Via Pessinetto 12, 10149 Torino, Italy
siragusa@di.unito.it

Ilias Tachmazidis
University of Huddersfield, Queensgate, Huddersfield, HD1 3DH, UK
i.tachmazidis@hud.ac.uk

Abstract

Knowledge representation and reasoning in the legal domain has primarily focused on case studies where knowledge and data can fit in main memory. However, this assumption no longer applies in the era of big data, where large amounts of data are generated daily. This paper discusses new opportunities and challenges that emerge in relation to reasoning with legal big data and the concepts of volume, velocity, variety and veracity. A four-layer legal big data framework is proposed to manage the complete lifecycle of legal big data from sourcing, processing and storage, to reasoning, analysis and consumption. Within each layer, a number of relevant future research directions are also identified, which can facilitate the realisation of knowledge-rich legal big data solutions.

1 Introduction

Since the emergence of computational knowledge representation and reasoning (KR), the domain of law has been a prime focus of attention as it is a rich domain full of explicit and implicit representation phenomena. From early logic programming approaches to elaborate logic-based mechanisms for dealing with, among others, notions of defeasibility, obligation and permission, the legal domain has been an inspiration for generations of KR researchers [5] [13].

Regardless of the particular approaches that have been applied in literature, most knowledge representation and reasoning research in the legal domain shares a common underlying assumption that related knowledge and data are relatively small in size. This, in turn, leads to the assumption that solutions can be executable by relying on standard main memory capacities. These assumptions have increasingly been challenged since the emergence of big data and the proliferation of data that become available through organisations, sensor networks and social media. The transformative effect of big data and analytics is evident in many fields, including law, where more and more data are available, as explained next.

Since legislation is at the basis of and regulates our everyday life and societies, many examples of big data must comply with, and are thus highly dependent on, specific norms. For instance, huge amounts of financial transactions must follow strict regulations, while complex food supply chains with myriads of sensor-based tracking data must comply with food regulations in various countries.

Organisations are feeling increasingly overwhelmed with the expanding set of legislation and case law available in recent years, as a consequence of several so-called black swan events, such as the Great Recession of 2008 and the COVID-19 pandemic, among others. Since 2008, the banking industry alone has received more than $300 billion in the form of legal sanctions from public institutions [17]. This has led to increased attention on regulative technology (RegTech) and, in particular, FinTech, i.e. RegTech applied to the financial domain.

As the law becomes more complex, conflicting and ever-changing, more advanced methodologies are required for analysing, representing and reasoning on legal knowledge [3]. An emerging question is whether and to what extent current legal informatics technologies can address KR-related needs within such scenarios, which can collectively be termed “legal big data”. As detailed in the following section, research has focused primarily on improving legal KR while dealing with relatively small subsets of legislation and relevant data [12], [20], with less effort put on examining how proposed approaches scale, when larger bodies of law and relevant data are considered.

More recent approaches that are capable of handling large amounts of legal data rely exclusively on machine learning [1], following the significant attention that machine learning has received in recent years in various fields, including the legal domain. While the ability of such techniques to discover hidden patterns in data, deriving new insights and making predictions is undeniable, we argue that the legal domain also requires a form of explainable artificial intelligence, aiming to increase scrutability and trust from end users, while also addressing liability concerns. Note that we distinguish explainability from the more narrow term of interpretability that is most often associated with machine learning models (for a disambiguation between the two terms the interested reader is referred to [6]). Explainability in this context must ensure that results reflect “correct” legal reasoning, referring back to specific pieces of legislation that justify them and avoiding faulty associations derived from data that would not be put forward by humans. Explainability can be significantly assisted by symbolic approaches, such as the logic-based approaches mentioned at the beginning of this section.

In this paper we propose ways to leverage the opportunities that the big data era opens with regard to reasoning in the legal domain by first identifying challenges that are introduced by legal big data and which state-of-the-art reasoning approaches may not be equipped to handle; these challenges relate to one of the well-known V’s of big data, namely volume, variety, velocity and veracity. We then propose a layered solution framework for knowledge-rich legal big data with explainability in mind as a step towards addressing the identified challenges, highlighting relevant directions for research at the confluence of AI, big data and law.

The remainder of this paper is organised as follows. Section 2 provides a concise overview of approaches in legal reasoning. Section 3 attempts to define legal big data by positioning the V's of big data in the context of legal data processing and presenting relevant challenges. To address these challenges, Section 4 proposes a layered framework for legal big data that covers the complete lifecycle from sourcing, storing and processing, to reasoning, analysing and consuming legal big data. Finally, Section 5 concludes and proposes directions for future research in the context of legal big data.

2 State of the Art

A brief summary of the various approaches that have been proposed for reasoning with legal information follows; a detailed account of such research, spanning more than four decades, is provided by Prakken and Sartor [21]. One of the first legal reasoning approaches relied on classical logic programming and was shown to be effective on self-contained legislation such as the British Nationality Act [23], with the largest case handled involving 50 pages of legislation encoded in 500 rules.

Following the advent of the Semantic Web and the introduction of the OWL family of languages, several research efforts focused on examining whether description logics are a suitable candidate for representing and reasoning about legislation. A prime example is HARNESS [26], which shows that well established sound and decidable description logic reasoners such as Pellet (https://github.com/stardog-union/pellet) can be exploited for normative reasoning, if, however, a significant compromise in terms of expressiveness is made. HARNESS was also tested on regulation excerpts consisting of no more than 3 paragraphs. On a slightly larger scale, [7] used reified I/O logic, a modern formalism based on reification [14] [22], to formalise the GDPR in 966 if-then rules (github.com/dapreco/daprecokb/tree/master/gdpr). It has been proved that the computational complexity of reified I/O logic is lower than deontic logics based on possible-world semantics [24].

A common issue that arises when using classical or description logics for legal KR is the fact that they are monotonic: logical consequences cannot be retracted, once entailed. This is in contrast to the nature of law, where legal consequences have to adapt in light of new evidence and conflicts between different regulations must be accounted for and resolved. Therefore, it is natural to employ non-monotonic logic for the purposes of normative reasoning. An example is a decision support system for regulations on adverse drug experiences using Defeasible Deontic Logic [16]. However, this system requires approximately 8 hours to process 3 million records which correspond to incidents from only the first quarter of 2014. Defeasible Deontic Logic and other approaches, including Answer Set Programming and argumentation are compared in [11] in terms of their expressiveness, inconsistency handling, reasoning support and complexity. Three use cases of increasing size are considered, with the largest involving 8 paragraphs of the United States US Food and Drug Administration (FDA) legislation.

Other approaches are more relevant to common law systems, such as case-based legal reasoning using CATO [2] or view legal reasoning as a process of argumentation, with opposing sides attempting to justify their own interpretation, such as Carneades [15]. These approaches, however, do not differ from the previously mentioned ones in terms of scale, in that they are focusing on legislation and data of relatively small size.

Few approaches, to the best of our knowledge, are capable of handling large amounts of legal data and all of them rely on machine learning and on the availability of high-quality labelled data [1]. As mentioned in the introduction, such approaches are lacking in terms of explainability of results.

3 Legal Big Data

Big data are usually characterised across four distinct dimensions: volume, velocity, variety and veracity, also known as the 4 V's. In this section we position the 4 V's in the context of legal data processing, helping define what legal big data entails and highlighting emerging challenges and providing examples for each case.

3.1 Volume

Volume is the main characteristic that makes data "big". In a legal context, large, interconnected legal corpora such as the EU legislation are likely sources of legal big data. Traditionally, legal reasoning was focused on encoding a set of given legislation and consequently processing each case separately in order to acquire a verdict. However, such reasoning was mainly based on in-memory approaches utilizing a single machine. Thus, processing large amounts of data would result in either memory shortage or long process times. As an example, the state-of-the art tool RuleRS [16] requires approximately 8 hours on a laptop to process 342 MB of data with more than three million forty thousand records of adverse events related to drugs. This corresponds to only a quarter of a year and refers to only a fraction of the overall regulation.

Enforcing legislation is another example where large volumes of legal data need to be processed. Millions of financial transactions are performed daily worldwide; illicit activity such as money laundering should be detected within these transactions. In addition, food supply chains are being increasingly automated through the use of the Internet of Things with large amounts of data being generated by tracking goods. In order to ensure quality, all imposed regulations throughout the supply chain need to be followed.

3.2 Velocity

Velocity is the frequency of incoming data that needs to be processed. In terms of legal data, velocity should be seen in terms of enforcing existing legislation over large amount of dynamic input. Such input is unlikely to be live transcripts from public hearings as it would be challenging to gather enough live information that could be considered as big data. In addition, fairly static use cases such as the FDA Adverse Event Reporting System and building applications should probably be disregarded in terms of their velocity since in the case of FDA data, each dataset is released every quarter of a year.

On the other hand, one can envisage applications that would require a careful investigation of the impact data velocity has. More specifically, financial transactions could potentially require real-time monitoring of day-to-day activity. Such functionality would depend on processing large amounts of transactions within seconds. Any illicit activity such as money laundering should be detected as soon as possible, with legal inspection of every financial transaction being an integral part of the process. Moreover, similar challenges apply to food supply chains where regulation compliance should be part of the overall process. This is essential for food supply chains that operate internationally, thus meeting both local and global regulations would improve quality and reliability.

3.3 Variety

Variety is a major challenge particularly in the legal context, since legal data sources are disparate in origin, format, organisation and structure. Any use case could require the combination of both structured and unstructured data coming from various sources. More specifically, financial transactions are expected to generate structured data where each field is well defined and machine processable. However, in the general case where data is unstructured, such as plain text describing any possible domain, identifying key concepts and transforming the input into a meaningful format could prove to be a challenge.

Data given to the system as input should first be transformed into a well defined and machine processable format (e.g. JSON, XML, RDF). Nonetheless, such transformation is not trivial since integrating data coming from disparate sources is a well known challenge. Legal processing sets an additional challenge which is the translation of all available data into a single format in order to acquire a uniform set of facts.

3.4 Veracity

Veracity refers to the trustworthiness of data, which may be affected by incompleteness or inconsistency. For instance, geodata relevant to building applications and property development efforts could in many cases lack important or contain outdated information. Data coming from the FDA Adverse Event Reporting System and food supply chains would also need to be curated in order to accurately reflect current world views. However, data that is inadvertently incomplete or inconsistent constitutes a lesser challenge compared to deliberately false statements, since such statements might be carefully inserted to the system by malevolent sources.

Any system operating over legal data would require specific procedures that could ensure the validity of available data as well as the trustworthiness of the source providing each dataset. Nonetheless, veracity should be seen as a dynamic process where information can be updated at any point, thus illustrating the need for dynamic evaluation of the legal implications of new data.

4 A Layered Framework for Legal Big Data

In this section, we propose to address the challenges brought on by legal big data through a four-layer framework that is a specialisation of the generic framework proposed by IBM for delivering big data solutions [19]. Figure 1 illustrates the four different layers that cover the complete lifecycle of legal big data, starting from legal big data sources and followed by three more layers responsible for processing and storage, analysis and consumption of legal big data. For each layer, we briefly describe its components, functionality and relevant directions for research.

4.1 Sources Layer

This layer acts as the foundation of the framework and encompasses all sources that can generate legal big data. We differentiate between general and context-specific sources. General legal big data sources include:

- Legislation documents (e.g. national law, EU law or international conventions)

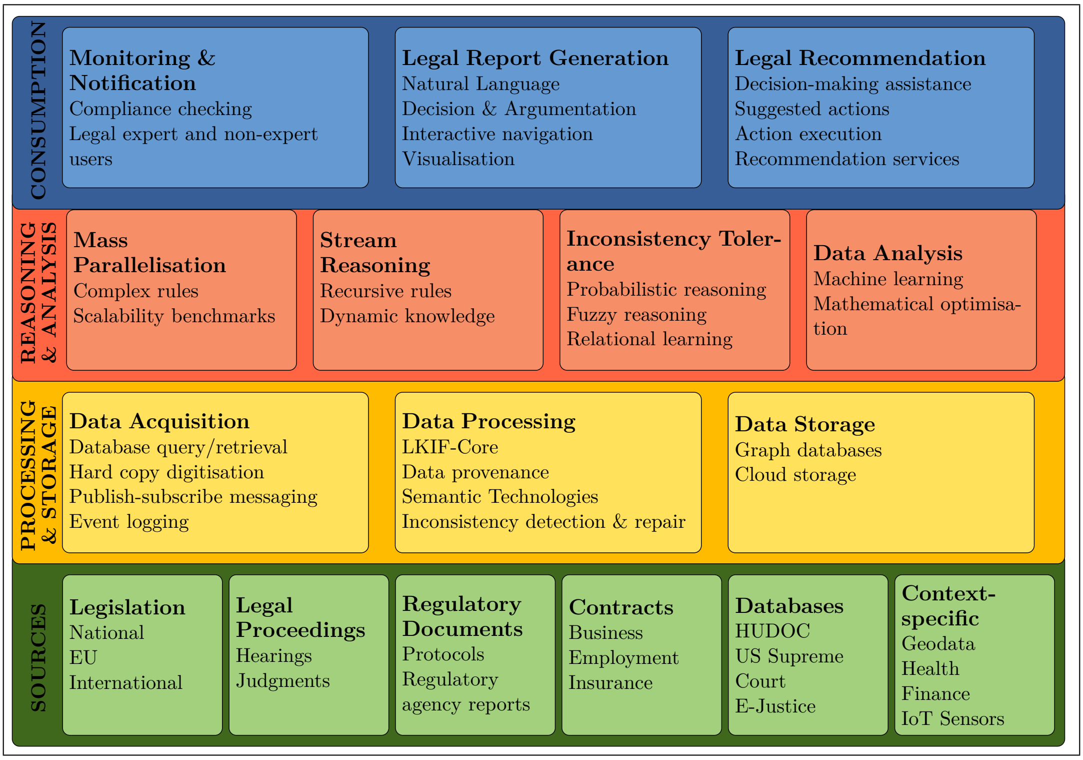

Figure 1: Overview of the proposed Legal Big Data layered framework.

- Legal proceedings (e.g. hearings or court judgments)
- Regulatory documents (e.g. protocols or regulatory agency reports)
- Contracts (e.g. business, employment, insurance)

Existing legal databases can also be used as sources for the framework, such as the European Court of Human Rights database (https://hudoc.echr.coe.int), the Supreme Court Database of the United States (http://scdb.wustl.edu) or databases available through the European E-Justice Portal (https://e-justice.europa.eu).

Apart from these data sources, legal applications may also indirectly depend on additional, context-specific sources. As discussed in the introduction legal big data may be generated from the following: geodata relevant to assess building and property/site development applications; health data related to drug regulatory approval processes; financial data that are monitored by legislation such as taxes; and IoT sensor data collecting within supply chains.

Given the vast differences between potential legal big data sources, the sources layer must be able to accommodate different setups across a number of dimensions. The first involves locating, accessing and collecting such data. Recent efforts in the context of open data initiatives such as the aforementioned databases can facilitate these tasks, providing central points of access where legal data can be accessed and searched. However, these efforts are still not widespread and cannot possibly cover any legal data requirement. Hence, in most cases some level of manual or ad-hoc data acquisition and aggregation method will be required, to ensure that all relevant data are available for the remaining layers of the framework.

With regard to the dimension of variety, data sources may come in many different formats, structured, semi-structured or unstructured ones. Legal proceedings, legislation, regulation and policy documents are almost always unstructured, in the form of simple, unannotated text. In exceptional cases, such documents may have been transformed into a semi-structured variant, using languages such as Akoma Ntoso or LegalRuleML (legalxml.org). In what concerns context-specific sources such as geodata or sensor data, structuredness is more likely but not guaranteed.

Finally, legal big data sources also vary in terms of rate of arrival. All general legal big data sources mentioned earlier generate new data at a very low rate, since drafting new legislation or trying a case in court, for instance, are time-consuming processes. High velocities can be witnessed in context-specific cases, such as financial and sensor data.

4.2 Processing and Storage Layer

This layer acts as a bridge between legal big data sources and the reasoning and analysis processes that will rely on them; hence its role is to prepare them so that they meet the requirements of these processes. This first involves acquiring legal big data, transforming/converting them to the required format, and then storing them.

Data acquisition for legal big data is as diverse as the sources described in the previous layer. In cases where data is already available in existing legal databases (or data stores), acquiring them means accessing these databases and retrieving all relevant data using appropriate queries. In other cases, a digitisation process may also be necessary to convert hard copies of legal documents to electronic counterparts. In what concerns context-specific sources whose rate of arrival is high, acquisition has to rely on streaming paradigms, such as publish-subscribe messaging or event logging.

Following acquisition, some form of preprocessing may be necessary. For instance, unstructured data cannot be used as-is for any reasoning or analysis task; they first need to be transformed to a semi-structured or structured form. Existing work on legal knowledge representation using semantic technologies can be leveraged for this transformation. Specialised legal ontologies have been proposed by researchers, such as the Legal Knowledge Interchange Format core ontology (https://github.com/RinkeHoekstra/lkif-core), which is capable of representing a wide range of legal concepts, from legal actions and roles to norms and argumentation. New representations may be required in case the concepts in LKIF-Core are not specific enough.

By semantically annotating legal big data based on ontologies, data structuredness is increased and a uniform representation is achieved across the different legal big data sources. This can also help in assessing data quality as well as identifying and handling inconsistencies. Research on data provenance and ontology repair can be applied for the case of legal big data, ensuring that reasoning and analysis results are not affected by unreliable or contradicting sources. However, such research has so far been limited to relatively small datasets [18] and research on large-scale data provenance and inconsistency detection and repair has not yet yielded any production-ready systems. Further research on parallelising diagnosis and repair over large-scale data could prove useful for legal big data as well.

Having ensured that acquired data are machine-processable, uniform and consistent, the penultimate step before storage is to perform any necessary format conversions to conform to the input requirements of the subsequent reasoning and analysis layer. This may be unnecessary in cases where representation and reasoning are closely related (e.g. RDF data and SPARQL reasoning). The data format also determines the most suitable storage solutions, such as graph databases or cloud storage.

4.3 Reasoning and Analysis Layer

This layer facilitates reasoning and analysis tasks in relation to legal big data by tapping onto research advances in relevant areas. One of these is mass parallelisation, which has been shown to be applicable to various types of reasoning [4]. Both supercomputers and distributed settings can be used in order to speed up data processing. The advantages are twofold, since mass parallelisation: (a) could significantly reduce processing time as multiple cores can be used simultaneously, and (b) virtually alleviates the restriction on main memory as more memory can be easily added to the system. However, there are certain issues that need to be addressed, as law would need to be encoded into some rule-based formalisation with potentially complex rules. In general, such rules tend to hinder mass parallelisation as novel optimisations and efficient rule evaluation techniques would need to be developed. In addition, legal big data itself (e.g. legal cases) would need to be studied in depth in order to comprehend the underlying patterns and data distribution, since data complexity might require special handling in order to ensure mass parallelisation.

In cases of high-velocity data, such as the context-specific data discussed in the sources layer, stream reasoning techniques may be harnessed. However, current state-of-the-art stream reasoning research [25] shows that only relatively simple rules could allow high throughput. In general, stream processing is intended for use cases where data is processed towards a single direction with recursive rules (i.e. rules that lead to inference loops) leading to performance bottlenecks. In addition, within such a dynamic environment, incoming data could potentially invalidate previously asserted knowledge leading to a new set of knowledge. Thus, reasoning on streaming legal big data depends on the development of novel techniques that can work with highly complex sets of rules.

Standard legal reasoning alone may not be enough to address use cases where legal big data face issues of uncertainty and imprecision, whether it is because of contradicting legislation or legal reports or because of inconsistencies in context-specific sources. As previously mentioned, data provenance and inconsistency detection methodologies may address some of these cases; if these are not successful, then inconsistency-tolerant reasoning approaches may be exploited. Some examples include probabilistic reasoning, which can consider situations where knowledge is associated with probabilities rather than truth values and fuzzy reasoning, deriving conclusions from approximate knowledge. Another alternative is the use of relational learning techniques, combining rules, probabilities and machine learning.

Apart from various forms of legal reasoning which facilitate explainability, this layer should also provide data analysis techniques based on data mining and machine learning. A prime example where such techniques can be useful is legal impact analysis. Assessing the impact of legal change requires determining optimal configurations out of a set of available alternatives or to simulate the impact of what-if scenarios. An interesting direction would be to investigate hybrid solutions that combine legal reasoning with well-established methodologies in simulation and optimisation. Considering, for example, the impact of traffic regulation on air pollution, data-driven models can be built to determine the short and long-term impact of specific traffic parameters, such as speed or volume. Legal reasoning could then pinpoint the specific regulation clauses that pertain to these factors. Alternatively, mathematical optimisation could also prove useful. For instance, to determine changes to traffic regulation that can keep air pollution within limits, legal reasoning can be performed on the existing regulation to determine traffic parameters, which can then be used to build a mathematical model representing their interaction and their effect on air pollution. Then, the goal would be to solve an optimisation problem that minimises the distance of air pollution levels from the legal limit.

4.4 Consumption Layer

The consumption layer is used for presenting the output of the above-mentioned layers, especially the reasoning and analysis layer, to the end users. These users can be legal experts or employees with or without specialized legal knowledge assigned with tasks involving legal reasoning such as monitoring legal compliance of actions, processes and policies in a company. For example a clerk sending information to third parties may receive a notification if a planned action violates data protection regulations. Thus the consumption layer is monitoring actions and notifies the end users accordingly.

In addition to monitoring and notification tasks, the consumption layer is responsible for producing reports for users that are typically not experts in computer science, but they often have a legal background. Reports must be produced in a human readable form, preferably in well-structured natural language typically used in legal documents. Reasoning results must be presented along with the reasoning steps that produced these results, allowing the end user to clearly understand (and be able to explain) the legal rules and preconditions, preferences and facts that yield specific results. The presentation must involve not only the final decision but also the different arguments used by parties involved in a legal dispute in order to evaluate the validity of these arguments, to further facilitate explainability. The consumption layer should offer the end users the capability to interactively navigate the justification output, since justification is equally important, if not more important than reasoning results. In order to increase readability, visualization techniques must be employed when producing reports. Therefore, legal design is another research direction worth investing in, working at the intersection among law, human-centered design, technology, and behavioural sciences.

Reporting functionality can be further extended by recommendation capabilities. In this case several scenarios are examined with respect to their legal consequences and based on the analysis results, the suggested actions are presented to the end users along with their justifications. The recommended action must ensure compliance with relevant legal norms and maximize several user-defined criteria. Since actions in the legal domain are the responsibility of humans, we expect the consumption layer to focus on deploying recommendation in the form of a service-based application [10, 9, 8] rather than involve automated execution of legal actions, with the possible exception of scenarios where legal tasks are repetitive and consequences of decisions are not critical. This is due to the fact that the legal domain is considered to be a critical part of social life and any decisions have to be scrutinised.

Finally, the consumption layer output may in turn be used as a legal big data source, creating a feedback loop with the lowest layer (or the second lowest, if the output is already stored in a machine readable and processable form). This could be especially useful in legal evaluation and reform procedures, offering insights on improving rules, policies and legislation in the future.

5 Conclusion

This paper argued that there is scope for new research in AI and law in cases where large amounts of legal knowledge and data are involved. We highlighted emerging challenges that are brought on by legal big data and positioned them across the dimensions of volume, velocity, variety and veracity, which involve: (a) handling large data volumes; (b) combining streaming data with existing legal knowledge; (c) integrating data from different sources and different formats; and (d) determining provenance and improving quality of available data. Examples where such legal big data may be produced range from the pharmaceutical, financial and property development sectors to food supply chains and impact analysis of regulatory change.

We then proposed a four-layer legal big data framework to address these challenges. The lowest layer concerns capturing the various legal big data sources, such as legislation, contracts and court judgments, as well as context-specific sources which are relevant to different legal use cases. The second layer is responsible for acquiring, processing and storing data coming from these sources. Various forms of reasoning and analysis processes are housed in the third layer and are tasked to produce results which are delivered to the end users through the top layer.

Through the description of the proposed framework, several directions for further research are identified. These are related to: (a) mass parallelisation of reasoning processes; (b) stream reasoning methodologies that can deliver high performance even within dynamic environments; (c) semantic technologies to represent, store and query large amounts of legal data; (d) data provenance, diagnosis and repair and reasoning approaches that can handle inconsistencies; and (e) hybrid solutions that combine simulation and mathematical optimisation with legal reasoning.

References

[1] Nikolaos Aletras, Dimitrios Tsarapatsanis, Daniel Preotiuc-Pietro, and Vasileios Lampos. Predicting judicial decisions of the European Court of Human Rights: a Natural Language Processing perspective. *PeerJ Computer Science*, 2:e93, 2016.

[2] Vincent Aleven. Using background knowledge in case-based legal reasoning:

A computational model and an intelligent learning environment. *Artif. Intell.*, 150(1-2):183–237, 2003.

[3] Grigoris Antoniou, George Baryannis, Sotiris Batsakis, Guido Governatori, Livio Robaldo, Giovanni Siragusa, and Ilias Tachmazidis. Legal Reasoning and Big Data: Opportunities and Challenges. In *MIREL 2018 Workshop on Mining and Reasoning with Legal Texts*, 2018. URL `https://doi.org/10.29007/tkmv`.

[4] Grigoris Antoniou, Sotiris Batsakis, Raghava Mutharaju, Jeff Z. Pan, Guilin Qi, Ilias Tachmazidis, Jacopo Urbani, and Zhangquan Zhou. A survey of large-scale reasoning on the web of data. *Knowledge Eng. Review*, 33:e21, 2018. doi: 10.1017/S0269888918000255.

[5] Grigoris Antoniou, George Baryannis, Sotiris Batsakis, Guido Governatori, Mohammad Badiul Islam, Qing Liu, Livio Robaldo, Giovanni Siragusa, and Ilias Tachmazidis. Large-Scale Legal Reasoning with Rules and Databases. *Journal of Applied Logics – IfCoLog Journal of Logics and their Applications*, 8(4): 911–939, 2021.

[6] Grigoris Antoniou, Emmanuel Papadakis, and George Baryannis. Mental Health Diagnosis: A Case for Explainable Artificial Intelligence. *International Journal on Artificial Intelligence Tools*, 31(3), 2022. doi: doi.org/10.1142/S0218213022410032.

[7] Cesare Bartolini, Andra Giurgiu, Gabriele Lenzini, and Livio Robaldo. Towards legal compliance by correlating standards and laws with a semi-automated methodology. In *BNCAI*, volume 765 of *Communications in Computer and Information Science*, pages 47–62. Springer, 2016.

[8] George Baryannis and Dimitris Plexousakis. WSSL: A Fluent Calculus-Based Language for Web Service Specifications. In Camille Salinesi, Moira C. Norrie, and Óscar Pastor, editors, *25th International Conference on Advanced Information Systems Engineering (CAiSE 2013)*, volume 7908 of *Lecture Notes in Computer Science*, pages 256–271. Springer Berlin Heidelberg, 2013. ISBN 978-3-642-38708-1. doi: 10.1007/978-3-642-38709-8_17.

[9] George Baryannis and Dimitris Plexousakis. Fluent Calculus-based Semantic Web Service Composition and Verification using WSSL. In A.R. Lomuscio et al., editors, *9th International Workshop on Semantic Web Enabled Software Engineering (SWESE2013), co-located with ICSOC 2013*, volume 8377 of *Lecture Notes in Computer Science*, pages 256–270. Springer International Publishing Switzerland, 2014. doi: 10.1007/978-3-319-06859-6_23.

[10] George Baryannis, Kyriakos Kritikos, and Dimitris Plexousakis. A specification-based QoS-aware design framework for service-based applications. *Service Oriented Computing and Applications*, 11(3):301–314, 2017. ISSN 1863-2394. doi: 10.1007/s11761-017-0210-4.

[11] Sotiris Batsakis, George Baryannis, Guido Governatori, Ilias Tachmazidis, and Grigoris Antoniou. Legal Representation and Reasoning in Practice: A Critical Comparison. In *Legal Knowledge and Information Systems - JURIX 2018: The Thirty-first Annual Conference, Groningen, The Netherlands, 12-14 December 2018.*, pages 31–40, 2018. URL `https://doi.org/10.3233/978-1-61499-935-5-31`.

[12] Guido Boella, Luigi Di Caro, Daniele Rispoli, and Livio Robaldo. A system for classifying multi-label text into eurovoc. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law*, ICAIL '13, pages 239–240, New York, NY, USA, 2013. ACM.

[13] Guido Boella, Luigi Di Caro, Llio Humphreys, Livio Robaldo, Piercarlo. Rossi, and Leendert van der Torre. Eunomos, a legal document and knowledge management system for the web to provide relevant, reliable and up-to-date information on the law. *Artificial Intelligence and Law*, 4, 2016.

[14] Andrew S. Gordon and Jerry R. Hobbs. *A formal theory of commonsense psychology, how people think people think.* Cambridge University Press, 2017.

[15] Thomas F. Gordon, Henry Prakken, and Douglas N. Walton. The Carneades model of argument and burden of proof. *Artificial Intelligence*, 171(10-15): 875–896, 2007. ISSN 0004-3702.

[16] Mohammad Badiul Islam and Guido Governatori. RuleRS: a rule-based architecture for decision support systems. *Artificial Intelligence and Law*, 2018. ISSN 1572-8382.

[17] Phillip Mann. RegTech: The Emergence of the Next Big Disruptor. *International Banker*, 25 October 2017. URL `https://internationalbanker.com/finance/regtech-emergence-next-big-disruptor`.

[18] Kodylan Moodley, Thomas Meyer, and Ivan José Varzinczak. Root justifications for ontology repair. In *International Conference on Web Reasoning and Rule Systems*, pages 275–280. Springer, 2011.

[19] Divakar Mysore, Shrikant Khupat, and Shweta Jain. Big data architecture and patterns, Part 3: Understanding the architectural layers of a big data solution. Technical report, IBM, Oct 2013. URL `https://www.ibm.com/developerworks/library/bd-archpatterns3/`.

[20] Rohan Nanda, Luigi Di Caro, Guido Boella, Hristo Konstantinov, Tenyo Tyankov, Daniel Traykov, Hristo Hristov, Francesco Costamagna, Llio Humphreys, Livio Robaldo, and Michele Romano. A unifying similarity measure for automated identification of national implementations of european union directives. In Jeroen Keppens and Guido Governatori, editors, *Proceedings of the 16th edition of the International Conference on Artificial Intelligence and Law, ICAIL 2017*. ACM, 2017.

[21] Henry Prakken and Giovanni Sartor. Law and logic: A review from an argumentation perspective. *Artif. Intell.*, 227:214–245, 2015.

[22] Livio Robaldo. Independent set readings and generalized quantifiers. *The Journal of Philosophical Logic*, 39(1):23–58, 2010.

[23] Marek J. Sergot, Fariba Sadri, Robert A. Kowalski, F. Kriwaczek, Peter Hammond, and H. T. Cory. The British Nationality Act as a Logic Program. *Commun. ACM*, 29(5):370–386, 1986.

[24] Xin Sun and Livio Robaldo. On the complexity of input/output logic. *The Journal of Applied Logic*, 25:69–88, 2017.

[25] Jacopo Urbani, Alessandro Margara, Ceriel J. H. Jacobs, Frank van Harmelen, and Henri E. Bal. DynamiTE: Parallel Materialization of Dynamic RDF Data. In Harith Alani, Lalana Kagal, Achille Fokoue, Paul T. Groth, Chris Biemann, Josiane Xavier Parreira, Lora Aroyo, Natasha F. Noy, Chris Welty, and Krzysztof Janowicz, editors, *The Semantic Web – ISWC 2013*, volume 8218 of *Lecture Notes in Computer Science*, pages 657–672. Springer, 2013. ISBN 978-3-642-41334-6.

[26] Saskia Van de Ven, Joost Breuker, Rinke Hoekstra, and Lars Wortel. Automated Legal Assessment in OWL 2. In Enrico Francesconi, Giovanni Sartor, and Daniela Tiscornia, editors, *JURIX*, volume 189 of *Frontiers in Artificial Intelligence and Applications*, pages 170–175. IOS Press, 2008. ISBN 978-1-58603-952-3.

 Received May 2022