

# A Hybrid Artificial Intelligence Methodology for Legal Analysis

Monica Palmirani, Salvatore Sapienza, Kevin Ashley\*

**ABSTRACT:** The following study introduces “Hybrid Artificial Intelligence Methodology for Legal Analysis” (HAIMLA). It consists of a six-step method to design, develop, validate and deploy artificial intelligence (AI) systems for legal analyses that are built on asynchronous unsupervised and supervised techniques applied to legal texts serialised in the Akoma Ntoso XML standard. HAIMLA methodology is drawn upon the existing literature and case studies in AI & Law and it is grounded on consolidated philosophical approaches. Taken together, this background inspires design requirements that constitute the essential pillars of HAIMLA and new directions for future refinements and implementations.

**KEYWORDS:** Legal analysis; methodology; machine learning; hybrid AI

**SUMMARY:** 1. Introduction – 2. Related Work – 3. Philosophical Foundations of HAIMLA – 4. Design properties of HAIMLA – 4.1. Interdisciplinary and Intradisciplinary – 4.2. Iterative and Interactive – 4.3. Knowable and Human-centric – 5. Steps of HAIMLA – 5.1. Legal Analysis and Hypothesis Definition – 5.2. Preparation – 5.3. Exploration – 5.4. Annotation and Training – 5.5. Validation – 5.6. Output – 6. HAIMLA in a thought experiment: the use case of Drug Dealing in the Italian case law – 7. Discussion – 8. Final remarks.

## 1. Introduction

Computational models of legal reasoning are meant to enable collaborative activities between humans and computers aimed at solving legal problems thanks to the development of legal applications.<sup>1</sup> Such computational models can take a constellation of forms, which range from systems intended to predict the output of a case, support legal information retrieval (IR), predict liability, and so forth.

When discussing Artificial Intelligence (AI) systems in the legal domain, Verheij<sup>2</sup> noted, in his address to the seventeenth International Conference on Artificial Intelligence and Law (ICAIL'19), that “[i]n AI as law, AI systems are to be thought of as *hybrid critical discussion systems* (emphasis added), where different hypothetical perspectives are constructed and evaluated until a good answer is found”.

---

\* Monica Palmirani: Professor at CIRSFD - ALMA AI, University of Bologna. Mail: [monica.palmirani@unibo.it](mailto:monica.palmirani@unibo.it); Salvatore Sapienza: Research assistant at CIRSFD - ALMA AI, University of Bologna. Mail: [salvatore.sapienza@unibo.it](mailto:salvatore.sapienza@unibo.it); Kevin Ashley: Professor at University of Pittsburgh, School of Law, USA. Mail: [ashley@pitt.edu](mailto:ashley@pitt.edu). This study has been conducted with the support of the European Commission funds within ERC HyperModeLex (Grant agreement ID: 101055185). The contribution of Salvatore Sapienza is supported by Programma Operativo Nazionale (PON) “Ricerca e Innovazione” 2014–2020 CCI2014IT16M2OP005, by the Italian Ministry of University and Research. The article was subject to a double-blind peer review process.

<sup>1</sup> K. ASHLEY, *Artificial intelligence and legal analytics: new tools for law practice in the digital age*. Cambridge, 2020, p.3

<sup>2</sup> B. VERHEIJ, *Artificial intelligence as law*, in *Artificial intelligence and law*, 28, 2, 2020, 181.



Palmirani introduced a debate on the use of non-symbolic AI that, based mostly on the probabilistic approach, could produce accurate results but without any meaningful relevance for legal studies.<sup>3</sup> In particular, she underlines two levels of problems: i) lack of robust legal hypotheses due to a lack of real integration of the computer science and legal studies methodologies; ii) some congenital problems in the use of non-symbolic techniques based only on the plain text. Some examples, reported also in her paper are the following:

- i) Paragraph/Sentence vs. Structure. Machine learning (ML), supervised or unsupervised, works at the paragraph or sentence level and may not take into account the document's structure. ML cannot semantically connect portions of the provisions (e.g., obligation-exception, duty-penalty).
- ii) Text vs. Context. ML often works without additional information about the context of the provision (e.g., jurisdiction, temporal parameters); this means ignoring elements that are key to the legal domain (e.g., derogations, that is, the partial repeals of laws, depend on certain conditionals, a clear example being sunset clauses).
- iii) Prediction vs. Relevance. ML works mostly by applying probabilistic techniques based on a data series (e.g., cases, decisions, fines issued by an authority, etc.), and if a trend becomes widespread in the legal system, it is likely to be repeated by the statistical model even if the underlying legislation has changed. For this reason, in the legal domain, it is also very important to consider the relevance (including temporal relevance) of the legal phenomenon being analysed (e.g., new legislation). This aspect should be included in the ML model using specialize techniques (e.g., assigning weights to events) that have already been adopted in some industrial sectors where recent data are more important than past data.
- iv) Internal vs. External content. ML does not consider normative and legal citations (normative cross-references) as qualified parts of a legal provision. For ML, a citation is just a sequence of characters. Depending on the use case, this may make it necessary to recall the portion of the text cited and inject it in the dataset.
- v) Static vs. Dynamic. The content linked up by way of normative citations changes as the legal system changes over time (e.g., art. 3 will not be the same forever). ML cannot understand this semantic aspect, and for this reason and for some use cases we need to integrate each normative citation with the corresponding point-in-time version of the text, that is, the version that was current at the time of the cited decision.

Inspired by these premises, this article investigates the possibility of establishing a methodology capable of dealing with the limitations of probabilistic approaches while leveraging on their flexibility. It introduces a Hybrid Artificial Intelligence Methodology for Legal Analysis (HAIMLA), i.e., a foundational method for the development of AI & Law systems. By proposing such a method, this study aims to find new answers to the question regarding how to combine AI methods in hybrid ways in the field of legal analysis. This is consistent with emerging hybrid approaches in computation – like

<sup>3</sup> M. PALMIRANI, *A Smart Legal Order for the Digital Era: A Hybrid AI and Dialogic Model*, in *Ragion Pratica*, 2, 2022.

neuro-symbolic artificial intelligence<sup>4</sup> – which are still relatively unknown to the domain of legal informatics.

A literature review on the notion of hybrid AI & Law systems reveals the lack of a methodology to be used in interdisciplinary research that encompasses consolidated trends and emerging techniques. Therefore, our goal is to provide the foundations for a method that is both based on decades of research in legal informatics and interconnected with legal philosophy. Furthermore, recent contributions in ethics, eXplainable AI (XAI), and law clearly identify the need to develop AI systems for legal analysis that are also resilient to novel legislative trends (e.g., the forthcoming Artificial Intelligence Act by the European Union), trustworthy and human-centred.

This study first discusses some case studies and the existing literature (Section 2). Then, it identifies how the philosophical foundations of AI & Law and emerging trends in AI ethics can contribute to shape a conceptual methodological framework (Section 3). Then, the essential properties of HAIMLA, including interdisciplinarity, iterativity, interactivity, human-centricity and explainability, are identified (Section 4). Then, the six steps of HAIMLA are explained and discussed (Section 5). Finally, following a discussion on the implications of HAIMLA method (Section 6), final remarks summarise the study and set directions for further research (Section 7).

## 2. Related work

The AI & Law literature is vast.<sup>5</sup> It aims to solve complex problems, including legal interpretation,<sup>6</sup> argument mining,<sup>7</sup> rule extraction,<sup>8</sup> while managing temporal aspects of legal documents.<sup>9</sup> To do so, computational approaches have been classified in two broad categories: legal expert systems, which contain representation of rules represented in a declarative language which specifies conditions and conclusions and legal text analysis, in which a legal information retrieval system is deployed to analyse legal *corpora*.<sup>10</sup>

<sup>4</sup> P. HIZLER, M.K. SARKER, *Neuro-symbolic artificial intelligence: The state of the art*, Amsterdam, 2022

<sup>5</sup> K. ASHLEY, *Artificial intelligence and legal analytics: new tools for law practice in the digital age*. cit.; T. BENCH-CAPON et al., *A history of AI and Law in 50 papers: 25 years of the international conference on AI and Law*, in *Artificial Intelligence and Law* 20, 3, 2020, 215.

<sup>6</sup> T. ATHAN, G. GOVERNATORI et al, *LegalRuleML: Design principles and foundations*, in W. FABER, A. PASCHKE, (eds.) *Reasoning Web International Summer School*, Cham, 2015, 151.

<sup>7</sup> D. LIGA, M. PALMIRANI, *Classifying argumentative stances of opposition using tree kernels*, in *Proceedings of the 2019 2nd International Conference on Algorithms, Computing and Artificial Intelligence*, 2019, 17; H. PRAKKEN, G. SARTOR, *Modelling reasoning with precedents in a formal dialogue game*, in *Judicial applications of artificial intelligence*, 1998, 127. H. PRAKKEN, G. SARTOR, *Law and logic: A review from an argumentation perspective*, in *Artificial intelligence*, 227, 2015, 214.

<sup>8</sup> A. WYNER, W. PETERS, *On rule extraction from regulations*, in *Legal Knowledge and Information Systems*, Amsterdam, 2011, 113–122.

<sup>9</sup> M. PALMIRANI, *Legislative change management with Akoma-Ntoso*, in G. SARTOR, M. PALMIRANI, E. FRANCESCONI, M. BIASIOTTI (eds.), *Legislative XML for the semantic Web*, Cham, 2011, 101; G. SARTOR et al., *Legislative XML for the semantic web: principles, models, standards for document management*, in *Springer Science & Business Media*, 4, 2011.

<sup>10</sup> K. ASHLEY, *Artificial intelligence and legal analytics: new tools for law practice in the digital age* cit.; K. ASHLEY, *Automatically Extracting Meaning from Legal Texts: Opportunities and Challenges*, in *Georgia State University*

These two approaches suffer from relevant drawbacks. A passage from Ashley<sup>11</sup> clarifies that:

“first, the techniques developed to enable expert systems to deal with uncertain and incomplete information tend to be *ad hoc* and unreliable; second, since the manual process of acquiring rules is cumbersome, time-consuming, and expensive, a knowledge acquisition bottleneck has limited the utility of expert systems in law; third, text analytics cannot solve this particular knowledge acquisition bottleneck. While emerging text analytics paradigms can extract certain kinds of semantic legal information from text, they are not yet able to extract expert systems rules.”

Within the field of text analysis, knowledge extraction and machine learning systems mirror the symbolic and connectionist approach to AI.<sup>12</sup> Knowledge extraction techniques require the explicit formulation of rules that correspond to patterns in the text (e.g., by regular expressions); instead, machine learning systems are capable of extracting patterns automatically.<sup>13</sup>

Therefore, a combination of these approaches in hybrid forms seems desirable as it might solve the knowledge acquisition bottleneck while maintaining some connection with expert rules. Scholars have attempted to develop connectionist and symbolic hybrid models. For instance, the “Symbolic and Connectionist Approach To Legal Information Retrieval” (SCALIR)<sup>14</sup> is an early attempt to combine the two approaches in an information retrieval system by a conceptual connection of network, weights, and thresholds that are determined by symbolic methods and progressively refined through users’ interaction. Similarly, LUIMA<sup>15</sup> uses machine learning approaches, namely naïve Bayes, logistic regression, and decision trees to detect patterns to be annotated by a knowledge extraction rule-based algorithm.<sup>16</sup>

Finally, two cases studies, respectively on privacy ontologies<sup>17</sup> and within the LEOS project (Drafting Legislation in the Era of AI and Digitization)<sup>18</sup> have shown the potential of hybrid AI methods to identify the nature of European Union corrigenda (i.e., error lists). Differently from the other examples,

---

*Law Review*, 35, 2018, 1117; L. ROBALDO et al., *Introduction for artificial intelligence and law: special issue “natural language processing for legal texts*. In *Legal Knowledge and Information Systems*, 2019.

<sup>11</sup> K. ASHLEY, *Artificial intelligence and legal analytics: new tools for law practice in the digital age*, cit., 11.

<sup>12</sup> S. RUSSELL, P. NORVIG, *Artificial intelligence: a modern approach*, London, 2010.

<sup>13</sup> K. ASHLEY, *Artificial intelligence and legal analytics: new tools for law practice in the digital age*, cit.; M. J. BOMMARITO, D. M. KATZ, E. M. DETTERMAN, *LexNLP: Natural language processing and information extraction for legal and regulatory texts*, in *Research Handbook on Big Data Law*, 2021. E. DE MAAT, R. WINKELS, *A next step towards automated modelling of sources of law*, in *Proceedings of the 12th International Conference on Artificial Intelligence and Law*, 2009, 266.

<sup>14</sup> D. E. ROSE, R. K. BELEW, *Legal information retrieval a hybrid approach*, in *Proceedings of the 2nd international conference on Artificial intelligence and law*, 1989, 138–146; D. E. ROSE, R. K. BELEW, *A connectionist and symbolic hybrid for improving legal research*, in *International Journal of Man-Machine Studies*, 35, 1, 1991, 1–33.

<sup>15</sup> K. ASHLEY, *Artificial intelligence and legal analytics: new tools for law practice in the digital age*, cit., 299.

<sup>16</sup> M. GRABMAIR et al., *Introducing LUIMA: an experiment in legal conceptual retrieval of vaccine injury decisions using a UIMA type system and tools*, in *Proceedings of the 15th international conference on artificial intelligence and law*, 2015, 69–78.

<sup>17</sup> M. PALMIRANI et al., *PrOnto ontology refinement through open knowledge extraction*, in *Legal Knowledge and Information Systems*, Amsterdam, 2019, 205–210.

<sup>18</sup> M. PALMIRANI et al., *Hybrid AI Framework for Legal Analysis of the EU Legislation Corrigenda*, in *Legal Knowledge and Information Systems*, 2021, 68–75.

another form of hybridisation of AI tools was adopted. In fact, the study made use of unsupervised and supervised machine learning techniques. Unsupervised learning was used for dataset exploration, whereas supervised experiments were conducted for automated labelling of legal provisions on the basis of a light-taxonomy progressively refined by linguistic signals extracted in the unsupervised phase. Incidentally, the legal text was also enriched by Akoma Ntoso XML mark-ups to preserve the legal semantics of the text. The conversion is performed to provide the texts with information that pertains to legal concepts, such as the entry into force of the document, intra-textual and extra-textual legal references, agents, and so on. Such schema-based conversion allowed a clearer explanation of the interaction between the extracted information and the legal text.

These samples show that two forms of hybridisation of AI tools are possible. On the one hand, some studies combined expert systems rules and machine learning; on the other hand, some studies relied on the latter computational approach and made use of unsupervised and supervised machine learning techniques. This was done asynchronously, with unsupervised learning being followed by supervised experiments. This paper will mainly discuss the combined use of unsupervised and supervised machine learning, yet bearing in mind the necessity of anchoring the methodology to the peculiarities of legal texts (e.g., the role of deontic operators, point-in-time references, etc.) that have been identified by the literature in legal expert systems. To do so, abstracting these experiments to a degree capable of identifying a general “hybrid AI” method should first investigate the philosophical foundations underpinning the examples shown above.

### 3. Philosophical Foundations of HAIMLA

Performing computational legal analysis is a complex task due to the nature of the legal text. Legal language conveys a meaning other than the simple signification of the natural text. Legal texts embed prescriptions, changes in the legal status of individuals, power attributions, normative references, values and principles, “open textures” and many other characteristics that make them different from other texts in natural language. While human beings, depending on their legal skills and background, are in most cases capable to understand the legal meaning of a regular text, computers do not share the same ability since they are relatively agnostic to the legal significance of the words under scrutiny.

Some scholars have advocated the impossibility, for machines, to “interpret a legal term or attribute a certain value to a legally protected interest” and “solve problems that follow from the semantic complexity of the law”.<sup>19</sup> This approach can be summarised in the formula “law is law, code is code”. Vice versa, some initiatives<sup>20</sup> are trying to “reverse engineer” legal reasoning practices by proposing a bottom-up, code-centred approach that bypasses the problems of legal interpretation by means of linguistic analyses.<sup>21</sup> This approach is usually referred to as the “law as code” movement.

---

<sup>19</sup> J. OSTER, *Code is code and law is law – the law of digitalization and the digitalization of law*. In *International Journal of Law and Information Technology*, 29, 2, 2021, 101–117.

<sup>20</sup> See the “Making Laws in a Post-Modern World: Are You Ready?” conference organised in September 2020 by the Canadian Institute for the Administration of Justice (CIAJ)

<sup>21</sup> See the preprint D. M. KATZ et al., *Natural Language Processing in the Legal Domain*, 2023. Available at <https://ssrn.com/abstract=4336224> or <http://dx.doi.org/10.2139/ssrn.4336224>



While the “law is code” approach might seem close to the hybrid methods discussed in this paper, the last thirty years of research in legal theory, interpretation theory, legal linguistics theory, and semiotics have identified solutions to some aspects of the semantic complexity of the law. Scholars have tried to bring the computational representation of legal texts (i.e., patterns, inferences, predictions, etc.) “closer” to the human understanding of legal documents by proposing several solutions. Neglecting their results and their underlying philosophy in a methodology-building process would lead to possible drawbacks in the results.

A method to tackle the issue of the computational representation of legal concepts is proposed by the Akoma Ntoso XML Standard<sup>22</sup> for the electronic representation of legal documents. By means of this standard, legal texts embed normative references that are meaningful to human beings and machine readable at the same time. Any computational operation performed on marked-up text is characterised by the semantic representation of the legal document that is consistent with the way in which legal operators understand the text. Then, isomorphism,<sup>23</sup> or the connection between the formal rules and the legally binding textual statements modelled by the rules,<sup>24</sup> ensures a “one-to-one correspondence between the rules in the formal model and the units of natural language text which express the rules in the original legal sources, such as sections of legislation”.<sup>25</sup> By means of isomorphic approaches – such as LegalRuleML – it is possible to formally represent the legal meaning of the text in a machine-readable way and, while doing so, preserve the human’s understanding of the rules embedded in the legal text. Similarly, legal ontologies (e.g., PrOnto<sup>26</sup>), used in combination with Ontology Design Patterns methods<sup>27</sup> allow a legally sound detection of entities – like agent, role, event, temporal parameters, action – in a way that mirrors the legal significance of the relationships within and among legal provisions.

As regards machine learning approaches, the promising results of the aforementioned studies may suggest that legal text analysis algorithms perform well in law-related tasks. However, the shift from symbolic to sub-symbolic computation calls for a different reasoning based on some properties of the latter paradigm. First, such representation of the legal text (i.e., patterns, inferences, predictions, etc.) is not determined by an *a priori* or pre-determined set of rules, but it is constructed by the machine. This construction might result from statistical correlations rather than a meaningful legal causation. The main issue raised by this factor is the necessity of a twofold validation under legal and technical points of view to ensure that results are valid from both perspectives. Second, results of

<sup>22</sup> M. PALMIRANI, F. VITALI, *Akoma-Ntoso for legal documents*, in G. SARTOR, M. PALMIRANI, E. FRANCESCONI, M. A. BIASIOTTI (eds.) *Legislative XML for the semantic Web*, Cham, 2011, 75–100.

<sup>23</sup> T. BENCH-CAPON, F. P. COENEN, *Isomorphism and legal knowledge based systems*, in *Artificial Intelligence and Law*, 1, 1, 1992, 65–86.

<sup>24</sup> T. ATHAN, *Oasis legalruleml*, in *Proceedings of the fourteenth international conference on artificial intelligence and law*, 2013, 3–12

<sup>25</sup> M. PALMIRANI, G. GOVERNATORI et al., *LegalRuleML: XML-based rules and norms*, in *International Workshop on Rules and Rule Markup Languages for the Semantic Web*, 2011, 298–312.

<sup>26</sup> M. PALMIRANI, M. MARTONI et al., *Pronto: Privacy ontology for legal reasoning*, in *International Conference on Electronic Government and the Information Systems Perspective*, 2018, 139–152.

<sup>27</sup> F. GANDON, G. GOVERNATORI, S. VILLATA, *Normative requirements as linked data*, in *JURIX 2017-The 30th international conference on Legal Knowledge and Information Systems*, 2017, 1–10; P. HITZLER, A. GANGEMI, K. JANOWICZ, *Ontology engineering with ontology design patterns: foundations and applications*, Amsterdam, 2016.



machine learning algorithms are represented by a probability rather than certainty, as they embed some degree of uncertainty due to the inner nature of abductive and Bayesian reasoning. While such margin of uncertainty might help dealing with the ambiguity of legal texts, validations from a legal point of view are necessary to correctly frame the degree of accuracy and the related metrics (True Positives, False Positives, True Negatives, False Negatives, Precision, Recall, F1). Moreover, a meaningful human control is deemed to be morally desirable in decision-making systems, according to the human-in-the-loop, human-on-the-loop, and human-in-command models.<sup>28</sup> Finally, some machine learning algorithms present significant challenges as regards the level of explainability.<sup>29</sup> Legal text analysis software requires an ability to explain conclusions<sup>30</sup> and due attention has been given to this issue in the AI & Law community.<sup>31</sup> A recent study<sup>32</sup> has pointed out the significant difference between the ways that have been used in existing AI-based legal systems to provide explanations (exemplary explanations for case-based reasoning; step-by-step explanations for rule-based reasoning; argumentation-based explanations; interactive explanations) and those used for machine learning systems in AI & Law. When discussing the latter, authors noted that “the rationales for the predictions (and hence the explanations) were often unsatisfactory”.

Most of the proposed solutions – if not all – are addressed to fulfil the necessity of reconciling the perception of the legal text by humans and computers. Human beings and machines build their own representations of the world when performing operations that require reasoning, and this includes reason with legal texts. Given the goal of reconciling human and machine perception of legal texts and the necessity of explainable outputs, results have to be sound both from a technical and a legal point of view. This entails that not only should computers perform well overall or in comparison to the state-of-the-art, but also their outputs should be legally meaningful, that is, significant for a legal operator beyond their statistical relevance. For instance, classifiers should be subject to an in-depth legal analysis aimed at understanding the legal meaning and implications of the results. Therefore, HAIMLA does not follow a “textualist” approach to legal analysis (“law is code”), but aims to include the broader meaning of legal texts into account by means of a responsible use of legal analytics tools that is trustworthy (e.g., in the selection of authoritative sources of law), explainable (i.e., in the meaning explained above), and accountable (i.e., with humans exercising control).

In sum, while HAIMLA is based on a “Hybrid AI” approach, it is also grounded on philosophical legal theory and axiological pillars. As regards the former, HAIMLA does not follow monolithic approaches, but places itself in-between the two extremes of “law is law, code is code” and “code is law”. In particular, it follows legal theories that support the computability of the legal language and the reconcil-

---

<sup>28</sup> See the document *Ethics guidelines for trustworthy AI* by the High Level Expert Group on AI set up by the EU Commission, 2019.

<sup>29</sup> See the seminal work of the AI4People consortium available at [https://www.eismd.eu/wp-content/uploads/2019/11/AI4People's-Ethical-Framework-for-a-Good-AI-Society\\_compressed.pdf](https://www.eismd.eu/wp-content/uploads/2019/11/AI4People's-Ethical-Framework-for-a-Good-AI-Society_compressed.pdf). See also the *Ethics guidelines for trustworthy AI*, cit.; T. MILLER, *Explanation in artificial intelligence: Insights from the social sciences*, in *Artificial intelligence*, 267, 2019, 1–38; F. PASQUALE, *The black box society*, Harvard, 2015.

<sup>30</sup> K. ATKINSON, T. BENCH-CAPON, D. BOLLEGALA, *Explanation in AI and law: Past, present and future*, in *Artificial Intelligence*, 289, 2020.

<sup>31</sup> See the “*EXplainable & Responsible AI in Law*” (XAILA) workshop at the 18th International Conference on Artificial Intelligence and Law (ICAAIL), 2021

<sup>32</sup> K. ATKINSON, T. BENCH-CAPON, D. BOLLEGALA, *Explanation in AI and law: Past, present and future*, cit.

Law

iation of humans' and computers' representation of legal texts. Moreover, it acknowledges the limitations of machine learning approaches and endorses the same values of the AI & Law community with regards to explainability and human oversight. To translate them into practice, HAIMLA encapsulates these philosophical foundations in design properties.

AI &amp;

#### 4. Design properties of HAIMLA

Design properties of HAIMLA aims to bridge the gap between the philosophical and axiological foundations discussed above and practical manifestations of the proposed methodology. These properties are meant to identify HAIMLA's essential components that shall be valid throughout each step and in the case of implementation by customisation. HAIMLA embeds these concepts by design, that is, from the early stages of the method.

##### 4.1. Interdisciplinary and Intradisciplinary

One core value of HAIMLA methodology is being interdisciplinary. Answering to AI & Law research questions, be they theoretical, explanatory through case studies, or fully applicative, requires the cross-fertilisation of at least two domains, AI and Law. This is consistent with the intermediary position between "law is law, code is code" and "law is code" approaches. As regards the technical domain, the field of AI disciplines is represented by the different approaches mentioned above. The field of legal studies should include *intra*-domain expertise depending on the object of the study (e.g., civil or criminal law) and *intra*-domain expertise (e.g., legal philosophy, computational law and data science). Cross-validation of the results from both perspectives is a crucial design property for the reasons given above.

##### 4.2. Iterative and Interactive

Iterations and interactions are crucial components of the collaborative relationship between computers and humans that HAIMLA conceptual framework intends to achieve. *Iterations* are cyclic steps that allow a progressive refinement of the computer's representation of reality with a human input. In turn, iterations allow a constant evaluation of any generated model or output by human beings. *Interactions* represent the two-way communication between technical and legal experts. In coordination with iterations, interactions foster the twofold spirit of the research initiative and the substantial cross-validation of the results from a technical and legal point of view. Therefore, HAIMLA proposes the inclusion of a legal analysis of sufficient depth to evaluate the report of metrics.

##### 4.3. Knowable and Human-centric

As with other domains, AI & Law follows the premises of XAI research. Therefore, the HAIMLA conceptual framework promotes explainable solutions both in the methodological approach and with regards to the output. From the methodological perspective, HAIMLA aims at integrating technical and legal expertise so that research results are constructed by a dialogic method and are understandable by both groups; from the perspective of the output, these results have to be interpretable,



explainable, and transparent.<sup>33</sup> In other words, the process and the output of any HAIMLA-generated output shall be knowable<sup>34</sup> to the widest possible audience, including the intended final users, the research community, the general public, and so forth. While doing so, HAIMLA promotes a human-centric approach and the development of beneficial AI systems in line with the axiological premises discussed above.

## 5. Steps of HAIMLA

HAIMLA (Fig. 1) comprises six steps, each in turn composed of tasks. Some of them (“legal tasks”) are attributed to a “legal team” made of experts in law, whereas others (“technical tasks”) are allocated to data scientists or individuals with similar expertise. Nothing prevents legal experts, however, from handling computational tasks, or vice versa, but it is crucial to separate the roles among the two teams to prevent invasive interactions during the research that might lead to “nudges” towards one field or the other. While HAIMLA is interdisciplinary, as it is addressed to mixed teams, a certain degree of separation is functional for independent building, evaluation, and validation of the final results. Some tasks addressed to identify the crucial aspects of the research (e.g., the research question(s), the hypotheses, the answer(s) to the question, placed at the beginning and at the end of the research process, are shared between the two teams. This is necessary to prevent imbalances in the task allocation and in the final results. Ideally, a team coordinator with mixed competences should coordinate the planning, allocate tasks to team members, supervise their execution and validate the outputs. To do so, each step shall be properly documented and traced.

### 5.1. Legal Analysis and Hypothesis Definition

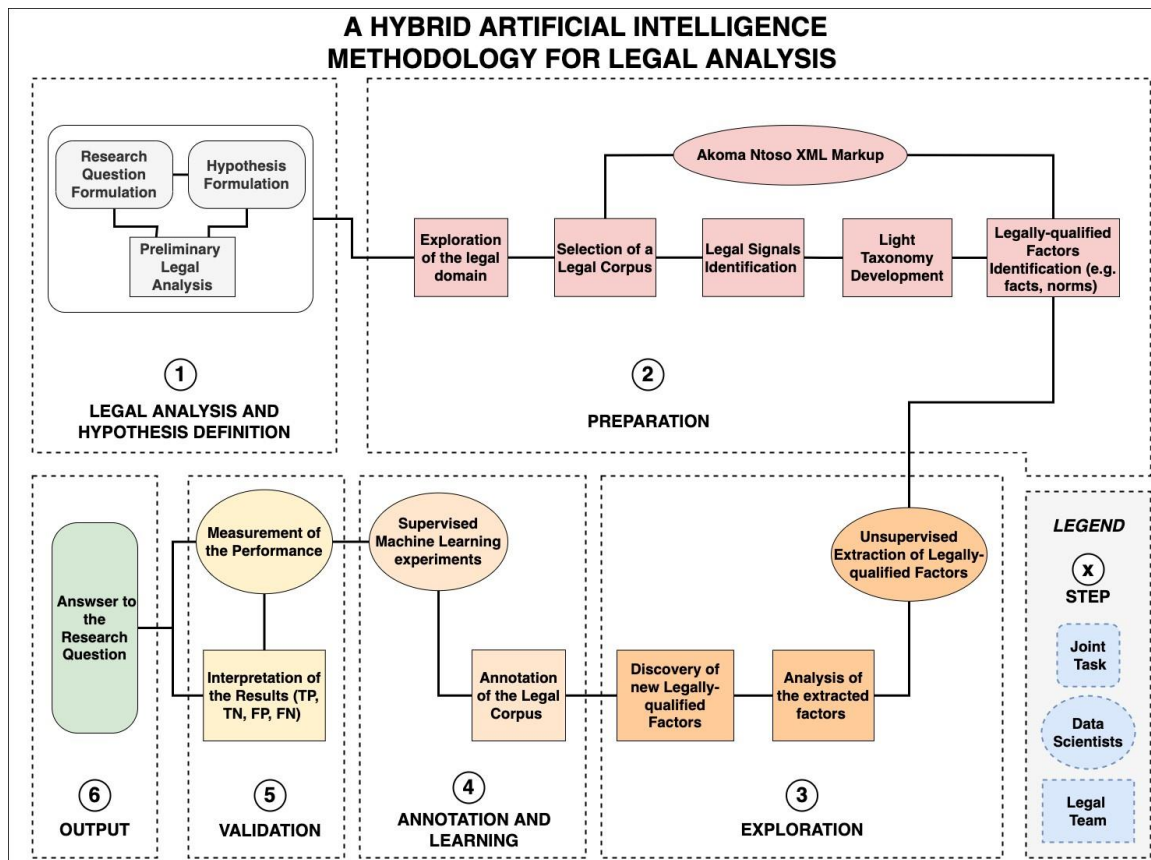
The “Legal Analysis and Hypothesis Definition” is devoted to planning the key aspects of the research activities. Both teams are involved in the process. These first steps consist of three sets of tasks:

- Research question(s) formulation. As with other domains, the identification of research questions to be answered is a crucial step in the definition of a legal analysis study. These questions can be identified both as pertaining to a legal sphere (e.g., “What is the relevance of the precedent *x* in the resolution of case *y*?”) or the technical one (e.g., “How does a legal Question-Answering system perform in retrieving a piece of legislation that matches the expectation of a human jurist?”)<sup>35</sup>. The question can

<sup>33</sup> A. BARREDO ARRIETA et al., *Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI*, in *Information fusion*, 58, 2020, 82–115. F. SOVRANO, S. SAPIENZA et al., *Metrics, Explainability and the European AI Act Proposal*, in *J – Multidisciplinary Scientific Journal*, 5, 1, 2022, 126–138.

<sup>34</sup> M. PALMIRANI, S. SAPIENZA, *Big Data, Explanations and Knowability Ragion pratica*, in *Ragion Pratica*, 2, 2021, 349–364.

<sup>35</sup> F. SOVRANO, M. PALMIRANI, B. DISTEFANO et al., *A dataset for evaluating legal question answering on private international law*, in *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law*, 2021, 230–234.



**Fig.1.** A Hybrid Artificial Intelligence Methodology for Legal Analysis (HAIMLA)

range from a more technical focus, especially when novel computational approaches are used, to a more legal one, when a tuned model is applied in a given legal domain. Literature reviews in AI & Law contribute to a better understanding of the current research trends, past examples and open questions.

- Hypothesis formulation. Closely related to the previous point, hypotheses are intuitive replies to the research questions, to be validated in the final step of the methodology. Technical and legal hypotheses should have an equal weight in the research process. As with research questions, a hypothesis can be formulated as technical (e.g., “method *x* performs better than method *y* in this set of tasks”) and legal (e.g., “case *a* is more relevant in the prediction of the outcome of case *b* in comparison to case *c*”).
- Preliminary legal analysis. This sub-task consists of broadly identifying the legal domain that will be analysed by the legal analytics tool. Legal formants method<sup>36</sup> can be used to identify and describe the essential applicable legislation, jurisprudence, and scientific literature on a given legal domain.

<sup>36</sup> R. SACCO, *Legal formants: a dynamic approach to comparative law (Installment I of II)*, in *The American Journal of Comparative Law*, 39, 1, 1991, 1–34.

## 5.2. Preparation

The “Preparation” phase aims to set up the conditions for the analysis. This preparation is meant both in technical (i.e., data pre-processing) and legal (e.g., legal research) terms. Some ordered tasks highlight the essential duties of both teams.

- Legal Tasks, ordered and parallel to technical tasks of this step.
  1. Exploration of the legal domain. Understanding in detail the legal domain is a crucial legal task. This helps to refine the research question and the hypothesis before entering into the operative parts of the project. In this stage, legislation, case law, and scientific literature are explored in depth.
  2. Selection of a legal corpus. On the basis of the (now refined) research question, a legal corpus (legislation, case law, and/or scholars’ opinion) is selected by the legal team. The size and the diversity of the corpus shall be appropriate to the research goals.
  3. Legal signals identification. The legal team is then in charge of identifying some recurring patterns in the legal text under scrutiny. A selection of the corpus is used in this step. For instance, in EU legislation, legal definitions are displayed in a form that is similar to the pattern “x” means “y”. The legal team should be able to identify patterns that signal potential legal relevance.
  4. Light taxonomy development. Having identified the recurring patterns in the legal texts, a light taxonomy of the domain is built. If necessary due to the complexity of the research question and if available, ontologies can be used.<sup>37</sup> For instance, since text classification can be binary, multi-class or multilabel, the taxonomy should match the intended scope of the analysis. To avoid excessive textualism, the built taxonomy should manifest a sufficient degree of abstraction that allows it to generalize most of the instances that can refer to the class. Legal categories from the literature are used to build the taxonomy. For instance, a taxonomy meant to allow the analysis of Italian criminal rulings should necessarily include, *inter alia*, *mens rea*, *actus rei*, and culpability as they constitute the pillars of Italian criminal law theory.
  5. Legally-qualified factors identification. The taxonomy is then expanded to include legally-qualified factors. The relationship between applicable norms, evidence, precedents, constitutional courts’ rulings, and the other factors discovered by the legal team should be identified to populate the taxonomy with some examples.<sup>38</sup> If the taxonomy is not sufficiently representative (i.e., too abstract or too textualist), it is refined.

---

<sup>37</sup> See the methodology in M. PALMIRANI, et al., *PrOnto ontology refinement through open knowledge extraction*, cit.

<sup>38</sup> See the domain model formalised in M. GRABMAIR, *Predicting trade secret case outcomes using argument schemes and learned quantitative value effect tradeoffs*, in *Proceedings of the 16th edition of the International Conference on Artificial Intelligence and Law*, 2017, 89–98.

- Technical tasks, parallel to legal tasks of this step.
  - XML Markup. In parallel with legal tasks, the 'technical team' marks up the legal text in Akoma Ntoso format.<sup>39</sup> This step is necessary to add a semantic and machine-readable structure to the text which carries the legal meaning of the textual references to norms and metadata that significantly enrich the capability of extracting legally-meaningful information from the text. If appropriate to the legal text under scrutiny, the LegalRuleML standard<sup>40</sup> can be used for legal rule modelling. This is crucial to maintain the isomorphic and defeasible connection between text and the rules embedded in the text, in a way that guarantees the provenance, authoritativeness, and authenticity of the norms. Moreover, LegalRuleML contributes to marking up deontic operators, parts of the norms, jurisdictions, and so forth in the text.

### 5.3. Exploration

The Exploration stage is meant to explore the legal corpus with the support of machines. In this step, unsupervised learning experiments are carried out to verify what information computers can extract automatically and how close their model of reality is to humans' understanding. At the same time, the outputs are used by the legal team to refine their model of knowledge, i.e., the taxonomy defined in the previous step.

- Technical tasks, prior to the legal tasks of this step.
  - Unsupervised Extraction of Legally-qualified factors. The technical team is tasked with the unsupervised extraction of legally-qualified factors, such as linguistic patterns or concepts that contribute to assign legal meaning to a given text. While this task might seem redundant with the previous one, it is necessary to extract knowledge as perceived by the computer. Without any supervision and with the least amount possible of information, the machine is asked to perform a classification task based on the taxonomy. Techniques such as Open Knowledge Extraction (OKE) or sentence/document clustering can be used in this step, also in combination with Knowledge Graphs.<sup>41</sup>
- Legal Tasks: ordered and following technical tasks of this step.
  1. Analysis of the extracted factors. Following the unsupervised extraction of legally-qualified concepts, the legal team analyses the results extracted by the algorithm and evaluates whether the taxonomy has a level of abstraction capable of distinguishing legal concepts/patterns while not being 'over-fitted' to the legal text under consideration.
  2. Discovery of new Legally-qualified factors: The legal team checks whether some factors, patterns, or concepts, found by the machine have escaped the taxonomy and new factors with legal relevance should be added to it. For instance, in one study the concept of 'minor' was discovered due to the automatic extraction of concepts.<sup>42</sup>

<sup>39</sup> M. PALMIRANI, F. VITALI, *Akoma-Ntoso for legal documents*, cit.

<sup>40</sup> M. PALMIRANI, G. GOVERNATORI et al., *LegalRuleML: XML-based rules and norms*, cit.

<sup>41</sup> M. PALMIRANI et al., *PrOnto ontology refinement through open knowledge extraction*, cit.

<sup>42</sup> M. PALMIRANI et al., *PrOnto ontology refinement through open knowledge extraction*, cit.

#### 5.4. Annotation and Training

“Annotation and Training” are the steps devoted to manually label the dataset and train the algorithms to perform a given task.

- Legal tasks, prior to the technical tasks of this step.
  - Annotation of the legal corpus. The aim of this step is to create a “gold standard” that is used to train, validate, and test the algorithm. To do so, legal experts perform the manual labelling of the legal document corpus according to the newly-defined taxonomy. The number of possible annotations depend on the nature of the task as defined in the first step (Section 5.1), as text classification can be binary, multi-class or multi-label. Manual annotation methods are widely deployed in AI & Law experiments.<sup>43</sup> A general consensus can be observed on the necessity of ensuring that annotators work independently, without suffering from interferences by external opinions, including and especially the ones of the other annotators or the technical team. Ideally, an odd number of annotators (e.g., 3) should be preferred because conflicts can be resolved by majority after with a conflict resolution phase led by an expert to train the annotators in this process. Finally, a law professor should review the annotations and certify them as a “gold standard”.
- Technical tasks, following the legal tasks of this step.
  - Supervised Machine Learning experiments. Once the “gold standard” is available, supervised machine learning experiments can be performed on the legal corpus. As with other cases, the literature in AI & Law is abundant, in particular in the field of classification tasks. The setup illustrated by Ashley<sup>44</sup> summarises the most relevant preprocessing sub-tasks, which include normalisation, stemming, stop-word erasing, tokenisation, and so forth. Following vectorisation, the dataset is partitioned in training, validation, and test sets. Percentages of the partitioning might range from 60-40 to 90-10 for training and test respectively, with the former set split into a “core” training set and an optional validation set with similar methods. Finally, one or more algorithms is trained, validated, and tested on the “gold standard”.

#### 5.5. Validation

The “Validation” step consists of the scrutiny of the results elaborated by the algorithm. This step ensures that a certain degree of human control is also exercised.

- Technical Tasks, prior and contextually related to the legal task of this step.
  - Measurement of the Performance. Results from classification algorithms are usually reported in confusion matrices that display True Positives, False Positives, True Negatives, False Nega-

---

<sup>43</sup> K. ASHLEY, *Artificial intelligence and legal analytics: new tools for law practice in the digital age*, cit., 301; M. GRABMAIR et al., *Introducing LUIMA: an experiment in legal conceptual retrieval of vaccine injury decisions using a UIMA type system and tools*, cit.; F. SOVRANO, M. PALMIRANI, F. VITALI, *Deep learning based multi-label text classification of UNGA resolutions*, in *Proceedings of the 13th international conference on theory and practice of electronic governance*, 2020, 686–695.

<sup>44</sup> K. ASHLEY, *Artificial intelligence and legal analytics: new tools for law practice in the digital age*, cit. chapter 8(3); F. SOVRANO, M. PALMIRANI, F. VITALI, *Deep learning based multi-label text classification of UNGA resolutions*, cit.

tives,<sup>45</sup> together with derived metrics, including Precision, Recall, and F1. First, the technical team records these metrics and transfers them to the legal team together with the relevant documentation. Then, it performs and independently evaluates of the results in comparison to the state of the art collected in the first step (Section 5.1). Furthermore, algorithmic biases should be assessed, in particular in the case of an unbalanced training dataset. To do so, classification results for each sub-class should also be made available to the legal team alongside the aggregated ones.

- Legal Tasks, following and contextually to the technical task of this step.
  - Interpretation of the result. Once the confusion matrix is received by the legal team, it performs a careful scrutiny of the performance of the classifiers from a legal perspective. In particular, False Positives and True Negatives shall be evaluated to understand what the limits of the classifier in categorising the legal text. With class-specific results, the legal team should evaluate the legal significance of sub-optimal performance of the algorithm with regards to specific legal elements (concepts, arguments, signals). The broader implications of biases and sub-optimal classifications of certain classes shall also be contextualised with regard to the possible deployments of the trained model.

## 5.6. Output

Finally, an answer to the research question is formulated and the hypothesis is confirmed or denied. The answer consists both of legal findings, intended as an assessment of the soundness of the algorithmic approach, and technical results, intended as overall algorithmic performance or in comparison to previous studies. Importantly, an assessment of the overall level of explainability and of possible biases is jointly carried out by the two teams. Research outputs are then drafted and made available to the public in various forms. These can include dashboards, visual data analysis tools, legal Q&A systems,<sup>46</sup> releases of the datasets and the codes created during the experiments, academic papers, and so on. Legal design techniques can help the visualisation of the results for final users without compromising their free interpretation of the results and visualisation-driven fallacies.

## 6. HAIMLA in a thought experiment: the use case of Drug Dealing in the Italian case law

To validate HAIMLA from a practical point of view, it is possible to hypothesize an application scenario of this method. In this regard, it is necessary to illustrate some legal peculiarities of the drug micro-trafficking (the trade of small quantities of illegal substances) an ideal scenario for its validation. When the criminal judge is called to evaluate certain facts and to carry out the process of syllogistic subsumption in relation to specific criminally relevant cases, the problem arises of ensuring maximum transparency in the path followed and, at the same time, uniformity and coherence of the *decisum* compared to the “line of precedents” that occurred on similar events. Despite Italy being a Civil

<sup>45</sup> K. ASHLEY, *Artificial intelligence and legal analytics: new tools for law practice in the digital age*, cit. chapter 8(3);

<sup>46</sup> F. SOVRANO, M. PALMIRANI, B. DISTEFANO et al., *A dataset for evaluating legal question answering on private international law*, cit.



Law country, the principles of predictability of the criminal sanction and formal equality implies that a certain degree of relevance is also given to the precedents. The affinity can be measured by the similarity between the factual elements present in the case brought to the attention of the judge and the previous ones and by the juridical proximity of the relevant institutes. To preserve equality in criminal justice, it would be necessary, in principle, to support the judge with AI systems that allow the rapid identification of similar precedents, in a factual and legal sense, to allow him or her to receive assistance in weighing up its decision also in terms of the quantum of sanctions.

The issue arises, indeed, with urgency with respect to those cases – such as in the hypotheses of micro-trafficking which are frequently brought to the attention of the judge. Due to the frequent verification and the easy identification of the typical elements (criminal conduct, object of the crime, criminally relevant quantity, subjective element), micro-trafficking can assume value also for what concerns the automated analysis of legal texts. Since these cases are characterized by a certain repetitiveness in the forms of manifestation, the factual and juridical elements that emerge from the cases of detention, sale, assignment, purchase of illegal substances constitute a useful test bench for testing the advantages and disadvantages that derive from combining the criminal judgment with innovative tools aimed at the automated analysis of sentences.

Most cases of micro-trafficking are characterized by the frequency that is technically necessary to hypothesize the semi-automated identification of constituent elements of the criminological type attributable to these criminal hypotheses directly on the text of the previous case law. Once identified, these elements can be further extracted, refined, and analysed, also by means of other legal data analytics tools, to facilitate the work of the judge who is faced with similar facts. Therefore, they constitute a useful benchmark for HAIMLA. It should be remarked that additional examples of similar hybrid methodologies can be found in other domains, including the activity of Audit Courts in Italy<sup>47</sup> and EU corrigenda.<sup>48</sup>

The table below represents how to perform an HAIMLA-backed analysis of cases of micro-trafficking. The represented scenario is the Italian Case Law, and the hypothetical research question is whether or not the quantity and quality of the illegal substance traded can be used to classify the crime as a misdemeanour or a felony.

---

<sup>47</sup> L. DECKER, et al. *Hybrid Classification of Audit Court Decisions using Online Context-Driven Neural Networks*, in Proceedings of the 2nd International Workshop on Knowledge Management and Process Mining for Law, 2023

<sup>48</sup> M. PALMIRANI, et. Al, *Hybrid AI framework for legal analysis of the EU legislation corrigenda*, in *Legal knowledge and information systems*, Amsterdam, 2022

Step	Legal Task	Legal Tasks: output	Technical Task	Technical Tasks: output
1. Legal Analysis and Hypothesis Definition	1.1. Research question formulation	Can the quantity and quality of the illegal substance traded in the case be used to classify the crime as a misdemeanour or a felony?	1.1. Research question formulation	What computational approach works best for such classification problem?
	1.2. Hypothesis formulation.	The seriousness of the crime depends on the quality and quantity of the traded substance. Therefore, these should be predictors for the classification.	1.2. Hypothesis formulation.	Algorithm "X" should perform better than the others
	1.3. Preliminary Legal Analysis	Recollection of the applicable law and of the relevant case law.		
2. Preparation	2.1 Exploration of the legal domain	Detailed analysis of the legal domain, which include the case law from high courts, constitutional courts, and legal doctrine.		
	2.2. Selection of a legal corpus	Selection of cases to be used as dataset. The composition of the dataset reflects the frequency of cases in local courts	Akoma Ntoso XML Markup	Once selected, the corpus of legal documents is converted to AKN Format
	2.3. Legal Signals Identification	Quantity is identified by a number followed by a unit of measure; Quality is represented by the closed list of		

		Table 1 annexed to DPR 309/1990		
	2.4. Light Taxonomy Development	A taxonomy of drug dealing is developed to connect the elements and the roles (e.g., “trader”, “buyer”), actions, penalties, sanctions, etc..		
	2.5. Legally-qualified Factors Identification	Linguistic formulae, including references to articles in law, are associated to the taxonomy. For instance, “Condanna ai sensi dell’art. 73 comma 5 del DPR 73/1990” in the decision is associated with misdemeanours, because the article applies drug dealing cases in which the combination of quantity and quality of the illegal substance is considered as a “minor offense”		
3. Exploration			Unsupervised extraction of legally-qualified factors	An attempt to classify the legally qualified factors extracted from the cases with the classes of the taxonomy with unsupervised classifiers. The classifier performs poorly on the matching

				of some drug-dealing actions.
	3.2. Analysis of the extracted factors	A legal analysis of mismatches. Drug-dealing actions (verbs) are misinterpreted by the classifier. Such actions have legal relevance since they should be included within the list of “drug-dealing actions” in DPR 309/1990 according to the case law.		
	3.3 Discovery of new Legally-qualified factors	A new list of drug-dealing actions is made, and the taxonomy is updated with the classes of actions that have legal relevance. New classes include “mediating”, “offer via social media”, and “exchanging for cryptocurrencies”.		
4. Annotation and Learning	Annotation of the Legal Corpus	Annotation of a the cases by an odd number of annotators, validated by a law professor (“gold standard”)		
			Supervised Machine Learning experiments	Identification and classification of the Legally-qualified factors according to the taxonomy. The

				algorithms can recognise quantity and quality of the drug, the relevant actions and the punishment
5. Validation	Interpretation of the Results (TP, TN, FP, FN)	False Positives often occur in the identification of 'drugs' that are expressed in chemical formulas or in slang (e.g., "speed" for "amphetamine")	Measurement of the Performance	A comparative table of the performance between the algorithms highlights the best classifier
6. Output	Answer to the Research Question		A research product discusses how the quantity and the quality of illegal substances can be used to automatically classify cases of drug dealing as felonies or misdemeanours.	

## 7. Discussion

The present study raises the possibility that a hybrid and cross-disciplinary methodology can be used in legal analysis. It combines supervised and unsupervised machine learning techniques applied to pre-processed, XML-enriched texts. The benefits of such an approach become clear when mixed teams - legal and technical - have to face the challenges of legal text analysis and are willing to adopt a cross-disciplinary approach. However, since HAIMLA is a rather conceptual framework, a full validation in practice is necessary to fully understand its potential and, possibly to refine it. To mitigate the lack of a practical counterpart, HAIMLA has been grounded on some experiments (reported in Section 2) and drawn upon a consolidated philosophical background (discussed in Section 3).

Adaptations of the HAIMLA methodology could become necessary in specific situations. However, the philosophy and the envisaged design properties of the proposed method should ease smooth adaptations of this conceptual framework to practical needs. In other words, HAIMLA's philosophical and design background constitute the "core" of the method, hence serving as a point of reference for future adaptations.

The HAIMLA methodology has also relevant implications for the development and deployment of legal "apps" that are compliant with the forthcoming AI regulations in Europe. In particular, observing the EU AI Act<sup>49</sup> is crucial to see the implications of adopting the HAIMLA methodology.

<sup>49</sup> European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (*Artificial Intelligence Act*) and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU)



AI systems that are “intended to be used by a judicial authority or on their behalf to assist a judicial authority in researching and interpreting facts and the law and in applying the law to a concrete set of facts, or to be used in a similar way in alternative dispute resolution” shall follow the stricter regime for “high risk” AI tools, as with in the thought experiment in Section 6.<sup>50</sup> Therefore, legal analysis tools that are meant to support judges or similar authorities shall comply with some requirements imposed by this legislation, as long as they qualify as AI systems.<sup>51</sup> Let us discuss how the HAIMLA methodology coordinates with some of these requirements. Article 10 (“Data and data governance”) of the Act states that training, validation and testing datasets shall be subject to certain practices, which concern, *inter alia* design choices, preparation stages, the formulation of the relevant assumptions, bias assessment and missing data strategies. HAIMLA allows a standardised and documented workflow – also relevant under Article 11 – that embeds these requirements *by design*. Article 13 (“Transparency and provision of information to deployers”) requires that high risk AI systems shall be designed in a way that their operation is transparent to the user and to the competent authorities. In particular, the characteristics, capabilities and limitations of performance of high-risk AI systems shall be made available to the user and, indirectly to the authorities (Annex IV). Two of the foundational pillars of HAIMLA are explainability and knowability. Hence attention must be given to metadata of the training, validation, and test datasets, the XML enrichment of these datasets, the documentation of each step, and the cross-validation of the performance. Article 14 (“Human Oversight”) requires that the design and development of AI systems shall be actively monitored by human beings. Design and development should allow human scrutiny of the AI system to an extent sufficient to correctly interpret its outputs. HAIMLA’s validation steps provide an opportunity for careful scrutiny of the results of the algorithm both from a technical and legal perspective. Finally, Article 15 (“Accuracy, robustness and cybersecurity”) paragraph 1, requires that high-risk systems shall be designed and developed in such a way that they achieve, in the light of their intended purpose, an appropriate level of accuracy. As with the previous point, HAIMLA’s validation step mandates a double verification of the accuracy level, not only from a technical perspective, but also as regards the legal soundness of the outputs.

## 8. Final Remarks

This paper has proposed a methodology for the development of hybrid AI tools for legal analysis. First, a recollection of the state of the art in AI & Law has reported some attempts to use hybrid AI computational approaches, in the twofold meaning of symbolic versus sub-symbolic reasoning and supervised versus unsupervised learning. This study aims to support the latter approach by providing a viable method to perform such analysis, without neglecting the contribution of symbolic and sub-

---

2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828.

<sup>50</sup> This is due to Annex III, which lists the high-risk systems referred to in Article 6(2).

<sup>51</sup> AI systems are defined by Article 3(1), which refers to any «machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments».



symbolic systems. To do so, it first identifies the core philosophical background and the essential design properties of such hybrid methodology and then proposes a workflow based on these pillars.

To fulfil the task and the necessity of reconciling humans' and computers' representation of legal texts, the proposed method consists of the asynchronous use of unsupervised and supervised experiments performed on legal documents. The analysis is carried out on XML-enriched legal *corpora* to preserve the essential qualities of legal texts (e.g., point-in-time references, deontic, references to case law, and other legally-qualified factors, as seen in Section 6) and on the basis of a "light" taxonomy that defines the legal classes to be identified. In this conceptual framework, a legal and a technical team, managed by a coordinator, cooperate to perform a cross-disciplinary analysis.

Five steps, namely 1) legal analysis and hypothesis definition, 2) preparation of the documents, 3) exploration, 4) annotation and learning, 5) validation of the results, and 6) output, have been proposed. Each step comprises legal and technical tasks, to be performed by experts in the respective fields. Joint tasks at the beginning and at the end of the workflow are provided to plan the research and to consolidate its outputs respectively.

Further research could refine the current version of HAIMLA methodology, perhaps in the light of new computational paradigms, including those identified by recent studies.<sup>52</sup> Eventually, future studies could adopt HAIMLA as their research methodology and provide valuable feedback to this contribution.

*IS  
&  
Law*

---

<sup>52</sup> K. D. ASHLEY, *Prospects for Legal Analytics: Some Approaches to Extracting More Meaning from Legal Texts*, in *University of Cincinnati Law Review*, 90, 4, 2022, 5.