



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

in cotutela con Universitat Autònoma de Barcelona

DOTTORATO DI RICERCA IN
LAW, SCIENCE AND TECHNOLOGY
Ciclo 37

Settore Concorsuale: 01/B1 - INFORMATICA

Settore Scientifico Disciplinare: INF/01 - INFORMATICA

THE ROLE OF LEGAL TECHNOLOGY IN THE HARMONISATION OF EU LAW

Presentata da: Davide Audrito

Coordinatore Dottorato

Antonino Rotolo

Supervisore

Luigi Di Caro

Cosupervisore

Pompeu Casanovas

Esame finale anno 2025

The Role of Legal Technology in the Harmonisation of EU Law

Contents

Acronyms	9
Introduction	14
1 Legal harmonisation in the European Union	20
1.1 Multilingualism and legal harmonisation in EU law	25
1.1.1 Legal and political drivers of multilingualism	26
1.1.2 Legal basis: EU treaties and Regulation No. 1/1958	28
1.2 Legal harmonisation and the EU legislative procedure	30
1.2.1 Ordinary legislative procedure	32
Legislative initiative	33
The codecision procedure	34
1.2.2 EU legislation and legal harmonisation: selected cases	36
1.2.3 EU law-making and linguistic challenges	39
Linguistic challenges	40
1.3 The meaning of EU law	43
1.3.1 Interpretation of EU law	45
Literal interpretation	47
Systematic or contextual interpretation	47

Teleological interpretation	48
Comparative interpretation	49
1.4 Digital tools for EU legal analysis	50
2 Similarity metrics for literal and comparative interpretation	54
2.1 Related work	55
2.2 Comparing European procedural safeguards in criminal proceedings: a case-study	56
2.2.1 The CrossJustice project	58
2.2.2 Manual annotation of directives	59
2.2.3 Overview of the legal corpus	60
2.3 Methodology	61
2.3.1 Legal text processing	61
2.3.2 Implementation of the similarity metric	62
2.4 Results	63
2.4.1 Representation of legal texts	64
2.4.2 Heat map visualisation	66
2.5 Conclusions	67
3 Network analysis for contextual interpretation and judicial precedents	68
3.1 Related work	69
3.1.1 Identifying and linking norms	70
3.1.2 Legal text analysis and machine learning	72
3.1.3 Legal citation networks	74
3.2 Network analysis in EU health law	77

3.2.1 Methodology	78
A corpus-based approach	78
Graph-based NLP approach	80
3.2.2 Initial results	82
Annotation guidelines	82
Annotation agreement	83
Grouping norms with card sorting	85
Graph analysis	85
Feature set	88
3.3 Co-occurrence networks of recitals and (sub-)articles in EU legislation	88
3.3.1 Methodological framework	90
Identification of norms and the annotation process	91
3.3.2 Results	92
Illustrative examples of the networks	93
3.3.3 Classification	94
Representation of legal text	95
Network analysis metrics	96
Experimental setup	96
3.3.4 The annotation process	97
Annotation results	97
Discussion of annotation results	99
Manual classification of norm groups	100
Card sorting results	101
3.3.5 Graph analysis	103

3.3.6	Classification outcome	105
	Inferring interrelationships	106
	Inferring link norm categories	106
	The impact of network features	107
3.4	Judicial citation networks: a CJEU case-study	108
3.4.1	Methodology	110
3.4.2	Results	112
3.5	Conclusions and future work	115
4	A lightweight ontology and a harmonised glossary for multilingual and analogical interpretation	124
4.1	Related work	125
4.2	An analogical lightweight ontology for EU criminal procedural rights in judicial cooperation	128
4.2.1	Tools and methodology	132
	Theoretical basis	133
	Ontological framework	138
4.2.2	Data collection and analysis	140
4.2.3	Ontology creation	145
4.2.4	Conclusions and future work	152
4.3	A harmonised glossary for LLM-based legal concept detection	154
4.3.1	Dataset	156
4.3.2	Methodology	159
	Manual concept identification	159
	Interconnecting concepts in a taxonomic structure	161
	A knowledge base for LLM-based concept-detection	163

Evaluating LLM-based concept detection	167
The implementing tool: Neo4J	168
4.3.3 Results	169
The EU concept-based glossary	169
LLM-based retrieval of national concepts	173
4.3.4 Conclusions	178
Conclusions	180
Appendix	185
4.4 System prompt	185
4.5 Prompt for legal concept detection	186
4.6 Annotation guidelines	189
4.6.1 Task Context	189
4.6.2 Task Objective	190
4.6.3 Annotation procedure	190
Input	190
Task 1	191
Task 2	192
Annotation Tool	192
Important Considerations	192
Bibliography	232

Abstract

The traditional doctrinal approach in legal studies, characterised by its rigor and reliance on precedent and formal reasoning, has long formed the bedrock of legal analysis. However, the increasing contribution of empirical and computational methods marks a significant evolution in legal scholarship. This dissertation explores the transformative impact of legal technology, particularly natural language processing (NLP) and knowledge modelling (KM), on the analysis of legal harmonisation within the European Union.

The contribution of this thesis lies in the integration of applied legal tech to examine EU legal harmonisation — a critical yet complex process aimed at aligning domestic laws under a common EU framework. From a legal-linguistic perspective, this first involves analysing the key constituents of the phenomenon, including multilingualism, the nature of language and legislative sources, and legal interpretation, all of which play a pivotal role in shaping the harmonisation landscape.

The core of the thesis, however, encompasses various hybrid legal technology methodologies for effectively analysing legal harmonisation within the European Union. The dissertation illustrates impactful applications through case studies, particularly in the context of EU-funded projects that enhance judicial cooperation across European Member States. By bridging legal and technological domains, the thesis contributes to understanding how digital tools, when implemented within robust methodological frameworks, can concretely enhance EU legal harmonisation, especially through the most recent generative AI technologies.

Moving from traditional vector-based computational pipelines to network analysis, ontological modelling, and language models-based approaches, this research proposes methodologies that, while accounting for the complexity of the phenomenon, can have a meaningful, concrete impact on the workload and capacity of European public administrations, with additional positive effects on legal practice. The outcome of this research,

conducted in the rigorous legal domain where results must meet the highest standards, is further validated by legal experts through manual knowledge modelling and their validation and discussion of the results.

The conclusions assess the challenges and prospects of deploying Large Language Models (LLMs) and related technologies in analysing legal harmonisation, while also evaluating research findings. These cutting-edge technologies hold promise for accelerating the comparison and exploration of multilingual legislative texts. Despite these advancements, incorporating accurate, manually-crafted data remains essential. This approach is crucial for minimizing the risk of producing misleading and inaccurate information.

Acronyms

AI: Artificial Intelligence

APIs: Application Programming Interfaces

avBetC: Average Betweenness Centrality

avCloC: Average Closeness Centrality

avClCo: Average Clustering Coefficient

avCons: Average Constraint

avDegC: Average Degree Centrality

avEigC: Average Eigenvector Centrality

avLoaC: Average Load Centrality

BERT: Bidirectional Encoder Representations from Transformers

BetC: Betweenness Centrality

BiLSTM: Bidirectional Long Short-Term Memory

BM25: Best Match 25

BoN: Bag-of-ngrams

CDM: Common Data Model

CELEX: Communitatis Europae Lex

CFLO: Current-flow Betweenness Centrality

CFREU: Charter of Fundamental Rights of the European Union

CJ: CrossJustice

CJEU: Court of Justice of the European Union

ClCo: Clustering Coefficient

Cliq: Clique

CloC: Closeness Centrality

CN: Co-occurrence network

Cns: Constitutive
Cnt: Contextual (meta-norm)
Conn: Connectivity
Cons: Constraint
CS: Cosine Similarity
De: Deontic
DegC: Degree Centrality
Degr: Degree
DG: Directorate-General
DL: Deep Learning
DT: Decision Tree
DTM: Document-Term Matrix
Du: Dummy rule-based classifier
EAEC: European Atomic Energy Community
EAW: European Arrest Warrant
ECC: European Economic Community
ECHR: European Convention on Human Rights
ECLI: European Case Law Identifier
ECSC: European Coal and Steel Community
EdBC: Edge Betweenness Centrality
EEA: European Economic Area
Effi: Efficiency
EigC: Eigenvector Centrality
ELISE: European Institutions Linguistic Information Storage and Exchange
ELTS: European Legal Taxonomy Syllabus

ET: Explicit Transpositions

EU: European Union

EURATOM: European Atomic Energy Community

FD EAW: Framework Decision European Arrest Warrant

FRBR: Functional Requirements for Bibliographic Records

GC: General Court

GDPR: General Data Protection Regulation

G1: Graph 1

G2: Graph 2

GU: Gazzetta Ufficiale

HG: Harmonised Glossary

HTML: HyperText Markup Language

IATE: InterActive Terminology for Europe

ICCPR: International Covenant on Civil and Political Rights

ICT: Information and Communication Technologies

Jacc: Jaccard Coefficient

JHA: Justice and Home Affairs

JSON: JavaScript Object Notation

kCli: k-clique

kNN: k-Nearest Neighbors

L4LOD: Licence for Linked Open Data

LDA: Latent Dirichlet Allocation

LDR: Linked Data Rights

LKIF: Legal Knowledge Interchange Format

LLMs: Large Language Models

LoaC: Load Centrality

LOIS: Lexical Ontologies for Legal Information Sharing

LR: Logistic Regression

MDS: Multidimensional Scaling

MEPs: Members of the European Parliament

ML: Machine Learning

NB: Naive Bayes

NGOs: Non-governmental organisations

NIMs: national implementing measures

NLP: Natural Language Processing

NLTK: Natural Language Toolkit

NM1: Network Metrics group 1

NM2: Network Metrics group 2

NM3: Network Metrics group 3

NMT: Neural Machine Translation

Ob: Objective

ODRL: Open Digital Rights Language

OEG: Ontology Engineering Group

OWL: Web Ontology Language

POS: Part-of-speech

Pr: Procedural (meta-norm)

RAG: Retrieval Augmented Generation

RDF: Resource Description Framework

RESTful: Representational State Transfer

RNNs: Recurrent Neural Networks

RSS: Really Simple Syndication

Sc: Scope

SMEs: Small and Medium-sized Enterprises

SVM: Support Vector Machines

TED: Tenders Electronic Daily

TFEU: Treaty on the Functioning of the European Union

TF-IDF: Term Frequency-Inverse Document Frequency

TEU: Treaty on European Union

TT: Transposing Text

URI: Uniform Resource Identifier

Weig: Weight

XML: eXtensible Markup Language

Introduction

Legal studies have traditionally been grounded in a doctrinal approach, widely recognised as rigorous and structured, founded on principles of precedent and formal reasoning. This approach combines case law, legislative and soft law analysis, while also incorporating various interpretative methodologies [1]. The emergence of quantitative legal analysis marked a pivotal shift towards adopting an empirical approach aimed at providing effective and measurable insights into legal phenomena [2]. This paradigm shift aimed to reduce the inherent arbitrariness often found in traditional legal analysis, where the discretion of expert analysts has a relevant impact.

Legal quantitative analysis in jurisprudence relies on empirical methods to potentially yield different outcomes compared to traditional methodological frameworks that allow legal professionals the arbitrary choice of case identification, which seems to define certain phenomena within jurisprudence. This approach has sought to validate legal theories through empirical evidence, though it has typically been constrained by the manual collection of data, limiting its scope quantitatively [3].

Transitioning from quantitative to computational analysis in law, which is still considered marginal despite recent technological advancements, marks the next evolutionary step [4, 5]. The advent of generative AI and natural language processing (NLP) technologies has highlighted the capabilities of software, such as chatbots, to handle complex legal arguments and autonomously perform routine legal tasks. One of the significant distinctions between quantitative and computational analysis lies in

the volume of data that can be managed. While quantitative legal analysis aims to validate legal theories with manually collected data, computational methodologies, leveraging big data, knowledge modelling techniques and NLP, allow for the automatic identification, extraction, and management of large amounts of legal documents, albeit not without limitations.

While this research does not primarily involve the handling of large datasets, the nature of the subject — specifically, the analysis of legal harmonisation within the European Union using computational methodologies — requires a focused and targeted approach. The scope of the study is tailored to address the challenges posed by legal harmonisation, a complex process aimed at unifying national laws under a common EU framework [6]. This process is intricate, involving multiple layers of legal, social, linguistic [7], and political [8] elements that are analysed to understand the broader implications of such harmonisation efforts.

This thesis explores the integration of hybrid legal technology designed to analyse the complex process of legal harmonisation within the European Union. The application of these methodologies is central to the research, with a particular focus on how they can address challenges faced in aligning national legal frameworks with EU directives and regulations. Drawing on advanced legal-tech approaches, the thesis examines a range of case studies developed in the context of EU-funded projects aimed at strengthening judicial cooperation across EU member states. These concrete applications demonstrate the benefits of incorporating technology into the legal sphere, highlighting how digital tools can facilitate more effective cross-border collaboration and harmonisation efforts.

More precisely, this research demonstrates their potential to make a substantial and measurable impact on the efficiency and capacity of European public administrations. These methodologies possibly aim to assist when parliamentary and governmental bodies are tasked with adopting implementing legislative measures of EU laws, in the case of member states, and when they need to assess the effectiveness and coherence of national laws in relation to EU legislation, in the case of European institutions.

The inability of EU member states to effectively transpose EU directives and regulations into their domestic legal systems has led to a significant number of infringement procedures initiated by the European Commission [9]. This issue underscores not only the complexities inherent in the transposition process but also the profound implications for regulatory and cultural integration across the EU. Beyond improving legislative drafting, the proposed methodologies also hold promise for legal practitioners, offering tools that could enhance legal interpretation and decision-making.

The research evolves from traditional vector-based computational pipelines to more sophisticated techniques, mostly including network analysis, ontological modelling, and LLM-based approaches. These methodologies are not only theoretically robust but also offer practical solutions that account for the multifaceted nature of legal harmonisation. The research outcomes have been rigorously tested within the stringent context of legal standards, where accuracy and reliability are paramount. Legal experts have played a key role in this validation process, manually modelling knowledge and assessing the results. Their feedback and validation lend additional credibility to the findings, confirming the practical relevance and applicability of the proposed methodologies in real-world settings, considering the sanctions that the European Union imposes to member states due to inconsistent domestic implementations of European legal acts.

This thesis is guided by one main research question and is structured across five chapters. The first chapter, which has a compilatory nature, aims to provide a broad overview of the core factors that influence and drive the harmonisation efforts, including the nature of language, linguistic variations, and judicial interpretations that play a crucial role in shaping legal uniformity within the EU. This introductory overview, succinctly covers the key drivers of legal harmonisation from a legal-linguistic perspective.

Although the methodologies employed are designed to study the phenomenon from a purely textual and semantic organisation perspective, as will be clarified in the next subsection, it has been deemed necessary to outline the most significant features of legal harmonisation in Europe. This

ensures that the phenomenon is not reduced to a simple textual analysis, thereby acknowledging the full complexity of the harmonisation process. A significant section of the first chapter has indeed been dedicated to the presentation of the interpretative methodologies most commonly employed by the Court of Justice of the European Union (CJEU), as hermeneutics plays a crucial role in legal harmonisation. Understanding how the CJEU interprets and applies EU law is essential to comprehending the broader dynamics of legal harmonisation, where interpretation shapes the practical implementation and alignment of laws across member states.

On top of the first chapter, the dissertation illustrates some of the legal tech methodologies employed throughout my Ph.D journey, which are guided by the following two research questions:

To what extent can legal technology support the analysis of legal harmonisation within the European Union, and can different legal-tech methodologies be integrated to achieve semi-automation of legislative harmonisation analysis?

In line with the research questions, the initial effort involves more traditional computational methodologies, with a case-study involving text mining techniques, including vector-based similarity metrics used to compare European legislative texts and their corresponding national implementations. On top of previous comparable approaches, the research path begins with a first attempt to validate the use of this methodology through a set of transposition tables, where the alignment between European and national laws was manually conducted within EU-funded projects. The experiment demonstrates the usefulness of this approach, using heat maps to visually represent the results. From a hermeneutic perspective, this initial step demonstrates how similarity metrics can serve as a valuable tool, particularly in supporting the literal interpretation of legislative texts, as well as comparative interpretation, especially in the context of multilingual analysis.

With regard to legal interpretation, the third chapter presents three

case studies on network analysis applied to legislative and judicial texts, showcasing approaches that are particularly useful for systematic or contextual interpretation of legal acts. This chapter illustrates how, when working with large datasets, it is possible to reveal the semantic connections between the recitals of a legislative act, which serve as its preamble, and the legally binding articles contained within. Such a method is valuable for comparing legislative drafting traditions and for showing how the interplay between preambles and articles can influence the interpretation of provisions, thereby indirectly affecting the effectiveness of harmonisation.

The third case study further demonstrates how case law can be systematised through the recursive citation of rulings, both within individual judgements and across cases. This facilitates a diachronic reading of judicial decisions, allowing for the identification and tracing of legal principles and their evolution over time. Additionally, this method offers a sectoral overview of case law and potentially enables the integration of legislative sources to compare their content with rulings from European and national courts, or even quasi-judicial authorities, especially in the context of European and international law.

Most prominently, chapter four proposes two methodological frameworks that incorporate knowledge modelling techniques and enables NLP applications. An important driver of legal harmonisation is indeed multilingualism, which poses unique challenges. Therefore, a key research outcome of this thesis — an analogical lightweight ontology of explicit and implicit definitions — specifically addresses potential discrepancies arising from translations into the twenty-four EU official languages. This approach mitigates the risk of divergent legal interpretations across different linguistic versions.

The second contribution of chapter four is a harmonised glossary developed in the framework of the EU-funded Facilex project. The glossary has proven effective in supporting the ability of advanced language models, i.e. Claude-3-Sonnet, to automatically infer domestic legislative concepts on top of a manually-curated selection of EU legal concepts. The integration

of legal accuracy and technological innovation offers a promising pathway for future legal harmonisation efforts, including potential applications available to public administrations.

The integration of LLM-based applications not only facilitates the semi-automation of concept extraction but also enables similarity analysis, thus pushing the boundaries of the current state-of-the-art, towards a concept-based similarity index. The use of the harmonised glossary within the prompting technique advances the field by improving the accuracy of text analysis and information retrieval, particularly in legal contexts where nuanced interpretation of concepts is crucial. The inclusion of synonym analysis into the prompt ensures that conceptually equivalent terms are appropriately recognised and retrieved, thereby offering a more comprehensive and meaningful similarity comparison.

Chapter 1

Legal harmonisation in the European Union

The establishment of today's European Union was primarily aimed at creating an international and regional legal framework that consistently enforces uniform and harmonised rules based on core shared principles such as the rule of law, solidarity, freedom, democracy, and equality [10]. The founding treaties, EU legislation, and the jurisprudence of the European Court of Justice have been pivotal to shaping a unified legal system that transcends national jurisdictions. To illustrate, "EU policies on the harmonization of national contract laws are predicated on fostering the advancement of Europe's internal market", which is indeed described as "one of the most significant achievements" of the EU, enabling "citizens and businesses to utilize their fundamental freedoms, including the free movement of goods and services" [11]. This move towards harmonisation and uniformity mostly extends to legal areas where the EU has legislative authority, including the internal market, consumer rights, the environment, labour, mobility, and migration.

While European law increasingly influences the majority of legislative areas, the effectiveness of legal harmonisation and uniformity among EU member states is contingent upon various factors which will be explored in the next sections. These factors include the general vagueness and

complexity of language, the EU's legislative procedures and sources, the multilingual nature of EU law, and judicial interpretation at both EU and national levels. The case law of the CJEU has developed legal principles that require member states to adhere to the Court's interpretations of EU law. However, it should be noted that these principles do not apply universally to all concepts within European law. To illustrate, given the complexities involved in delineating the notion of family within EU law, Article 2, Paragraph 2 of Directive 2004/38/EC introduces a distinct framework for its definition. According to this directive, a family member encompasses not only the spouse but also the partner in a registered partnership with an EU citizen, assuming the member state affirms the partnership's equivalence to marriage. Furthermore, it includes direct descendants who are either under the age of 21 or dependants, as well as those of the spouse or partner, in addition to dependent direct relatives in the ascending line and those of the spouse or partner. The definition of family as per this directive is rooted in both biological and legal connections among its members [12]. Moreover, the inclusion of a provision for the "recognition of partnerships" highlights that the EU lacks the authority to uniformly define the family concept across its member states [13]. In conclusion, referring to the complexity of the phenomenon of legal harmonisation, some scholars argue that complete linguistic and legal harmonisation is an unattainable goal [14].

Given the characteristics of European legislative sources, directives as legislative instruments inherently tend to increase differences among EU member states even when they are implemented comprehensively. Accordingly, Article 288(3) of the Treaty on the Functioning of the European Union (TFEU) stipulates that a directive is binding on each member state to which it is addressed in terms of the result to be achieved, but it allows national authorities the autonomy to determine the form and methods of implementation. This mechanism of directives, employed by the European Parliament and the Council of the European Union, is aimed at harmonising national legal frameworks by mandating conformity in outcomes while

granting discretion in procedural aspects. Moreover, directives can be categorized into minimum and maximum harmonisation directives, which further influence the degree of legal convergence [15]. Minimum harmonisation directives set a baseline standard that all member states must meet but allow them to adopt more stringent measures if desired. This can result in divergent national laws even after the directive is implemented [16]. Conversely, maximum harmonisation directives aim to fully align national laws by setting both the minimum and maximum standards, thereby limiting the ability of member states to introduce additional or stricter provisions [17].

In contrast, EU regulations establish uniform rules that are directly applicable across all member states without requiring additional transposing legislation. The distinction between these two legislative sources has become less pronounced over time. This trend is supported by scholarly analysis, including a discussion on the General Data Protection Regulation (GDPR) [18], which essentially necessitated corresponding national legislation. The implementation of the GDPR resulted in varied national legislative frameworks that are often characterised by differing requirements. To illustrate, Legislative Decree No. 2018/101 entered into force in Italy, introducing measures to align Italian national legislation with the provisions of the GDPR. In addition to incorporating the stipulations of the GDPR, Legislative Decree No. 2018/101 also established regulations on certain matters within the purview of national legislative authority, including the specification of certain types of criminal offenses, alongside the monetary penalties already outlined by the GDPR. The example of the GDPR clarifies that the different methodologies proposed in this thesis are equally applicable to both directives and regulations. Although EU regulations do not require transposition through the adoption of national implementing legislation in member states, they often necessitate significant amendments to national legislative frameworks. Therefore, the methodology detailed in the final chapter is particularly valuable, as it facilitates the alignment of European legislation with national legislation.

By utilizing a standardized approach to legislative documents that allows for tracking amendments, this methodology would streamline the implementation process, ensuring consistency between regulations and national laws.

Despite the transposition of directives into national law ostensibly achieving the goal of establishing common rules, the pursuit of harmonisation faces significant challenges, as mentioned previously, due to the inherent flexibility afforded by directives. First, the transposition of EU legal concepts into the legal systems of member states is a complex process due to the diverse interpretations applied to the same legal texts across twenty-seven distinct political systems with their own traditions, and definitions [19]. A notable example is the term “spouse” as used in the EU Citizens’ Directive (Directive 2004/38/EC), which lacks a specific definition despite the CJEU having interpreted it on previous occasions [20, 21]. This omission poses a constitutional challenge, compelling the CJEU to mediate between conflicting interests. While eleven EU member states recognise same-sex marriages, seven define marriage as exclusively being a union between a man and a woman. Most legal systems acknowledge same-sex couples and accept civil partnerships legally established in another EU country. Furthermore, the interpretation of “spouse” is critical for the applicability of free movement law *ratione personae*, particularly when same-sex married couples move to a member state that acknowledges their marital status and thereby enjoy the protection of EU law. The significance is accentuated for couples that include a third-country national, who may only possess a derived right to free movement [22].

Second, the translation of EU legislation into twenty-four official languages exacerbates the complexities of interpretation, occasionally leading to the adoption of definitions that diverge from the original ones. At the EU level, under the principle of linguistic equality, divergences among different linguistic versions of an EU act cannot constrain the CJEU to adopting a single linguistic interpretation [23]. Instead, the court must consider the intent and overall framework of the rules in question with

so-called contextual or systematic interpretation, which is aimed at interpreting the applicable rule within the context of the the relevant legislative framework. This principle might even necessitate that one linguistic version should be disregarded in favour of others. Moreover, the CJEU is allowed to refer to particular linguistic versions to bolster its legal reasoning. For example, in *Stauder* [24], the CJEU was faced with an European Commission decision that butter should be offered at reduced prices to socially assisted consumers. The interpretation challenge arose because in two linguistic versions, the discount was conditional upon disclosure of the consumer’s name, while other versions required a ‘coupon referring to the person concerned’, which allowed alternative verification methods. A narrow interpretation could have undermined the uniform application of the decision. Therefore, the Court opted for the ‘most liberal’ interpretation [23], concluding that the Commission’s decision should be understood as not requiring — but also not prohibiting — the identification of beneficiaries by name.

In domestic courts, judges must apply the directive translated into their national language. In cases of interpretative uncertainty, judges are expected to conduct a comparative analysis of the various linguistic versions of the EU legal provision before deciding to apply the *acte claire* doctrine. This doctrine allows national courts to skip referring questions to the Court of Justice of the EU if the application of EU law is obvious and leaves no doubt. This principle, first established in the *CILFIT v Ministry of Health* case [25], streamlines legal processes by enabling national courts to resolve clear legal issues without needing guidance from the CJEU. The doctrine asserts that if EU law’s interpretation is certain and consistent across member states, and there’s sufficient CJEU precedent, a referral is unnecessary [26]. It is also crucial to recognise that the norms transposed into domestic legal frameworks may themselves be affected by the linguistic ambiguities inherent in EU legal provisions.

The subsequent sections in this chapter will offer a succinct overview of the primary factors driving legal harmonisation within the European

Union. These include multilingualism, legislative procedures and other legislative sources, as well as the interpretation of EU law by both the Court of Justice of the European Union (CJEU) and domestic courts.

1.1 Multilingualism and legal harmonisation in EU law

The multilingual character of European law significantly impacts the process of legal harmonisation across the legal systems of EU member states through EU legal acts. To this purpose, Theodor Schilling has observed that “no two texts in different languages will ever have exactly the same meaning” and further noted that “even significant divergences between different language versions of a text cannot fully be avoided”.

In examining the European Union (EU) as a unique case study within the broader spectrum of international organisations and multilingual states, it becomes apparent that the EU stands out due to its extensive multilingual regime [27, 28]. All 24 official languages of its member states are also recognised as official languages at the EU level. This setup is critical for enabling democratic deliberation and ensuring that EU law is equally authentic across all these languages, a necessity since EU legislation directly affects its citizens in various countries.

The decision to embrace robust multilingualism is driven by the desire to ensure equal treatment for all member states and to maintain unity within the EU’s diverse legal and cultural contexts [29, 30]. Despite the logistical challenges this poses in law-making and communication, the principle of “equal authenticity” remains a cornerstone of EU law [31]. This principle implies that all language versions of legal texts carry equal legal weight and has reached a “quasi-constitutional” status [7], pursuant to Article 55 TEU. Accordingly, the treaty, drawn up in a single original in the Bulgarian, Croatian, Czech, Danish, Dutch, English, Estonian, Finnish, French, German, Greek, Hungarian, Irish, Italian, Latvian, Lithuanian,

Maltese, Polish, Portuguese, Romanian, Slovak, Slovenian, Spanish and Swedish languages, the texts in each of these languages are equally authentic. More than just a legal requirement, the provision reflects a profound commitment to the linguistic diversity found within the EU, which includes over 80 languages spoken by its population [32].

While the EU manages this diversity by adopting institutional multilingualism, which is the EU policy on the use of languages within its institutions, it remains a sensitive political and legal issue, often subject to debate and calls for more streamlined approaches [33]. However, changing the EU’s linguistic policy requires unanimous consent from the Council, underscoring the complexity and sensitivity involved. The discussion also highlights the specialist skills required for multilingual legal drafting within EU institutions, with broader implications for the harmonisation and unification of laws across Europe [34]. The intricate relationship between law and language in the EU’s legislative process not only presents significant challenges but also demonstrates the Union’s distinctive method of governance and legislation in a multilingual setting.

1.1.1 Legal and political drivers of multilingualism

This section briefly summarises the legal and political reasons underpinning multilingualism in EU law and is aimed at introducing its legal basis arising from the TFEU and Regulation No. 1/1958 [7]. More precisely, Article 42 of the Charter of Fundamental Rights of the European Union (CFREU) lays down the right of access to documents, which qualifies as a fundamental right, whereas Regulation No. 1049/2001 on public access to documents is based upon the principle of the ‘widest possible access’ to EU documents [35].

Firstly, the European Union’s legal framework mandates multilingualism as a fundamental aspect of its legislative processes in order to ensure that EU law is accessible and comprehensible to all European citizens. This principle is rooted in democratic values such as legal certainty, judi-

cial protection and non-discrimination, which are crucial for maintaining the rule of law within the Union [36, 37].

Multilingual publication of EU legislation in the Official Journal is essential for the law to have direct effects in member states [38], and it encompasses, under specific conditions, EU directives [39]. This requirement stems from the principles of legal certainty and equality before the law, as established in the EU treaties and reinforced by case law. For instance, the Court of Justice of the European Union (CJEU) has emphasised the necessity of multilingualism for the effective application of EU law, particularly for ensuring that legal texts are equally authentic in all official languages [40]. The principle of equal authenticity prevents any linguistic version of legislative text from being considered subordinate, thus safeguarding the rights of individuals to rely on the version in their own language [33, 41].

Furthermore, the multilingual requirement supports democratic accountability and transparency. It ensures that European citizens can access and understand legislative documents, thereby facilitating informed participation in the democratic processes of the EU [42]. The Treaty of Amsterdam (now Article 15 of the TFEU) and Regulation No. 1049/2001 on public access to European Parliament, Council, and Commission documents underscore this, granting all EU citizens the right to access documents in their own language, which is vital for the full exercise of their democratic rights [43].

Secondly, from a political perspective, the commitment to multilingualism reflects the EU's foundational principles of respect for linguistic diversity and equality among member states [44]. This policy not only recognises the cultural significance of languages but also aligns with the political ethos of treating all member states equally, irrespective of their size or the prevalence of their languages [37].

The policy of linguistic equality is evident in the EU's institutional practice in which the official languages of all the member states are granted

equal status in EU proceedings. This approach reflects a broader commitment to cultural diversity, as stated in the Laeken Declaration, which emphasises respect for linguistic and cultural differences as central to the European identity [45]. The declaration views multilingualism as integral to sustaining the unique cultural identities within the EU, promoting a sense of inclusion and equal participation among all member states [46].

Additionally, multilingualism in legislative procedures has been defended on the grounds of political equality and transparency. The European Parliament has explicitly recognised these aspects as crucial for maintaining a democratic and transparent governance framework within the EU. By ensuring that legislation is drafted and published in all official languages, the EU facilitates equal access to legal texts, thereby supporting fair and democratic decision-making processes [47].

In summary, the EU's commitment to multilingualism is underpinned by a complex interplay of legal and political motives. Legally, it is driven by the principles of democracy, legal certainty, and the direct effect of EU law. Politically, it reflects the Union's dedication to linguistic diversity, cultural identity, and equality among its member states. These dual foundations not only reinforce one other but also contribute to the cohesion and integration of the EU, aligning its legal practices with its overarching political goals and values.

1.1.2 Legal basis: EU treaties and Regulation No. 1/1958

In the legal framework of the European Union (EU), the principle of full multilingualism, established by the founding treaties, underscores the equality of all official languages, precluding the dominance of any single language. This principle has remained unchallenged, forming a foundational aspect of the operations of EU institutions despite their increasing complexity due to the Union's expansion and the growing number of official languages — from 4 to 24 in just over five decades [48].

The EU approach to multilingualism is built upon the recognition of each member state's official language, as designated by their national constitutions. Consequently, the legal basis for the EU's linguistic regime has remained largely unchanged, apart from adjustments necessitated by the accession of new member states. This situation, while fostering democratic participation, introduces operational challenges as the Union now comprises 27 members, following the withdrawal of the United Kingdom [49].

The principle of equal authenticity, integral to EU multilingualism, ensures that all official languages have equal standing, reflecting the Union's commitment to preventing any single language or culture from becoming dominant [43]. The treaties and secondary legislation affirm that all texts are equally authentic in all 24 official languages. However, they do not extensively cover the implications of multilingualism for legal interpretation, leaving significant discretion to the European Court of Justice [50].

The primary legal instruments — the treaties, followed by Council Regulation No. 1/1958 — dictate the legal framework for language use within the EU, differentiating between treaty languages, official languages, and working languages [51]. The category of treaty languages, a formal, legal, language class, concerns the linguistic versions of the EU founding treaties that are considered authentic. There were four of them initially at the time of the Treaty of Paris in 1951 and this increased to twenty-four with the last accessions. When a country joins the European Union, its selected official language must be one recognised by its constitution. The European Parliament notes that there is no policy allowing member states to choose which languages to designate as their official EU language besides the one recognised by the domestic constitutional system [52]. This criterion has faced criticism for sometimes favouring languages spoken by fewer people, such as Irish, which became an EU official language in 2007 despite having only about 40,000 to 80,000 native speakers. Conversely, languages like Catalan, with over 10 million speakers, are not recognised as official EU languages despite their large speaker base [7].

Some years later, Regulation No. 1/1958 specifically established the

multilingual nature of the European Economic Community (ECC), mandating that legislation would be drafted and published in the official languages. The Regulation established the framework for official and working languages within the European Union. According to Articles 4 and 5, all regulatory texts and documents of general application must be drafted and published in the official languages, with each version being equally valid. Article 6 allows EU institutions to decide which languages to use in specific scenarios through their procedural rules, and this flexibility has proven to be significant in practice. Currently, there is no formal legal distinction between official and working languages [53]. However, in practice, working languages are typically used for internal communications within the institutions, while official languages are used for external communications, including the publication of legislation [54, 55]. The linguistic regime’s complexities are further highlighted by the ongoing adjustments required to fully include all languages in EU operations, exemplified by the gradual incorporation of Irish as a full official and working language. Overall, the EU’s steadfast adherence to multilingualism, despite the inherent challenges, underscores its commitment to equality and cultural diversity, serving as a critical component of its legal and democratic framework.

1.2 Legal harmonisation and the EU legislative procedure

The legislative procedure of the European Union (EU) is intricately designed to support the goal of legal-linguistic harmonisation across the member states’ legal systems. This procedure addresses the unique challenge posed by the EU’s commitment to maintaining linguistic diversity, requiring that legislative texts be available in all twenty-four official languages. The necessity of drafting legislation in multiple languages complicates the law-making process but is essential for ensuring accessibility and inclusivity [56].

Instead of co-drafting texts simultaneously in all languages — a method rendered impractical by the number of languages involved — the EU has developed a sophisticated mixed system where the processes of drafting, translating, and legal-linguistic revising are closely interlinked and alternate continuously [57]. This system begins with the drafting of a text in one language, followed by multiple cycles of translation, revision, and legal-linguistic checks. Each cycle is aimed at refining the texts to ensure that that translations are not mere linguistic conversions but are legally accurate and contextually appropriate across different legal systems [58].

This process challenges the notion of an “original” text [59] because each language version is crafted to stand as an authentic and authoritative text in its own right. The iterative nature of this drafting process is crucial for maintaining the legal quality and clarity of legislative documents, thereby preventing discrepancies that could hinder the uniform application of these laws [60].

By adopting this approach, the EU effectively promotes legal harmonisation, facilitating consistent interpretation and enforcement of laws across member states. This is particularly important in a union where legal and cultural diversity must be bridged to create a coherent and functioning supranational legal order. The legislative procedure, therefore, not only respects linguistic diversity but also enhances the legal integration necessary for the successful operation of the EU. This notwithstanding, and will be illustrated in the second part of this thesis, the structure of the ordinary legislative procedure does not entirely remove obstacles that hinder full, comprehensive legal harmonisation in the EU member states.

For the purpose of this thesis, the focus will be solely on: the ordinary legislative procedure through which the vast majority of legally binding acts are adopted, and the relationship between the complexity of language and the creation of legal norms.

1.2.1 Ordinary legislative procedure

The Lisbon Treaty classifies EU legal acts into three primary types — legislative, delegated, and implementing — establishing a clear hierarchy between legislative and non-legislative (executive) acts [61]. Despite the prevalence of non-legislative acts predominantly formulated and ratified by the Commission [62], the ordinary legislative procedure, once known as codecision, stands as the standard method for inter-institutional law-making. Following the adoption of the TFEU, the majority of the EU’s legislative actions are expected to be passed collaboratively by the European Parliament and the Council, following the legal framework set out in Articles 293 to 299 of the TFEU. In the legislative period from 2014 to 2019, nearly 90% of legislative proposals by the Commission were handled via this procedure [63].

The drafting of an ordinary legislative act is a concerted effort by the Commission, which submits the proposal, and the European Parliament and the Council, which are responsible for the act’s adoption. The actual drafting process does not constitute traditional joint drafting or merely translating from one language into 22 others. Instead, it involves a cyclical process of drafting, translation, and legal-linguistic revision, which begins at the Commission and continues through the Parliament and the Council [58]. This procedure is underpinned by a sequential framework, as described in the TFEU, whereby an institution’s decision regarding the text prompts a response from the other institutions stating whether to accept, reject, or modify it. The formal final vote is typically preceded by negotiations to reach a consensus on the text, a practice formalised by the “Joint declaration on practical arrangements for the codecision procedure” [64]. This negotiation process often takes place in trilogue meetings, where compromises are discussed in closed sessions, which raises concerns about transparency [65].

Linguistic considerations are also critical to this process. The drafting often starts with one language version as the base, which is not neces-

sarily maintained throughout the entire process. The proposal’s source language is primarily used for drafting during various procedural stages, including trilogue negotiations and discussions within Council working groups. Translations and legal-linguistic revisions are conducted once the substance of the draft has been finalised by all the institutions involved. The initial text may need to be retroactively adjusted if translations into other languages reveal errors or ambiguities in the original version [57].

The following section outlines the meticulous and collaborative nature of the EU legislative process under the Lisbon Treaty, highlighting the complexities of drafting, decision-making, and linguistic adjustments made by the Commission, Parliament, and Council. For this purpose, two short subsections address the legislative initiative and the codecision phase of the legislative procedure. Integrating these elements ensures the effective creation and adoption of legislative acts within the EU’s legal framework.

Legislative initiative

The ordinary legislative procedure in the EU, as per Article 293 of the TFEU, commences with the Commission submitting a proposal drafted by the relevant Directorate-General (DG) typically in English, the common drafting language despite most drafters being non-native speakers [58]. The initial drafting takes place within the middle administrative tiers of the competent DG and the General Secretariat of the Commission [66], and English has prevailed as the drafting language over the past decade, despite historical use of the French language [67].

The drafting process involves non-legal technical experts producing a preliminary draft, which is then refined through numerous revisions by other Commission staff [68]. This phase often reveals a reliance on linguistic precedents that may not suit new circumstances and could perpetuate past errors [67]. To address linguistic deficiencies, especially given that a significant majority of drafters are not native English speakers [69], the Commission’s DG Translation provides an editing service to enhance

linguistic quality, although this is not mandatory and is in fact under-utilised [68].

Once the DG has finalised the draft, it undergoes inter-service consultation, where the Legal Service plays a critical role in reviewing both legal and linguistic aspects of the draft in its source language [70]. Following this, the finalised text is translated into all official EU languages by the Directorate General for Translation (DGT), typically after approval by the College of Commissioners but before submission to the legislative bodies [58].

Upon submission, the Commission’s influence over the draft’s wording diminishes and constitutes primarily facilitating trilogues between the Parliament and the Council to reconcile differences and potentially alter the original proposal based on parliamentary amendments until the Council intervenes [58]. The presence of Commission drafters in trilogues is crucial to maintaining continuity and ensuring the drafters’ intentions are considered in discussions [71]. This structured yet flexible drafting process highlights the intricate balance between linguistic precision, legal scrutiny, and inter-institutional collaboration within the EU’s legislative framework.

The codecision procedure

The legislative drafting and approval process within the European Commission often reflects similar challenges to those of the European Parliament and the Council. A multitude of interveners from twenty-eight member states and 750 Members of the European Parliament (MEPs), many of whom are not specialists, must interpret and amend legislative texts that are predominantly drafted in English — a foreign language for most [72]. This process sometimes makes use of the French version during negotiations, though that is less common [68]. A significant portion of legislative acts, over 70%, are adopted at the first reading in accordance with the provisions of Article 294(3)–(6) of the TFEU. Pursuant to Article 294(7)–(9), if disagreements persist after the first reading, the

process advances to a second reading and potentially a conciliation procedure if consensus remains elusive, following the legal framework set forth in Article 294(10)–(12) of the TFEU.

During the codecision procedure, the drafting process involves multiple stages of review and amendment without a clear ‘master version’ of the legislative proposal, complicating linguistic consistency and legal clarity. At the Council, versions in languages other than English may influence amendments during the working group stages. However, the drafting and amendment process relies heavily on the use of English, with provisions for linguistic remarks and possible reservations being made early on in the process [58].

In the meantime, the European Parliament appoints a file coordinator responsible for overseeing the linguistic and legal quality of texts throughout the legislative procedure. This role is crucial given the Parliament’s multilingual nature, which necessitates extensive legal-linguistic review of all proposed amendments before and after committee approvals [73]. The inter-institutional translation and legal-linguistic revision process is carefully timed to ensure the text is ready for final approval, with the Parliament and Council collaborating closely to finalise the text [74]. This collaboration continues up to the plenary vote, where texts must be legally and linguistically solidified in all official languages to reflect the political agreements reached [73].

Ultimately, the complex interplay between drafting, negotiating, and translating highlights the intricacies of formulating legislation in the EU. The process ensures that all legislative texts are thoroughly vetted and linguistically aligned across all official languages, thereby maintaining legal integrity and facilitating legislative consensus across diverse linguistic and legal frameworks [68]. This rigorous process is designed to uphold the legislative quality and multilingual equity essential to the EU’s legislative integrity.

1.2.2 EU legislation and legal harmonisation: selected cases

This section deviates from earlier discussions by succinctly examining how the phenomenon of legal harmonisation is shaped not only by the processes of the ordinary legislative procedure and its inherent multilingualism but also by the characteristics of European legislative instruments, particularly directives and regulations. This analysis is supported by some examples that vividly illustrate these dynamics.

Early in 1959, Polach maintained that “the European Economic Community (ECC), Euratom (EAEC), and Coal and Steel Community (ECSC) has provided a new framework and impetus for the harmonization of laws in Western Europe” [75]. As mentioned in the introduction, directives require member states to achieve a designated outcome while granting them the autonomy to determine their methods of implementation. This flexibility allows for adaptation to the diverse legal systems and cultures within the European Union, thus facilitating a harmonised legal framework that still respects national variations. In other words, legislative harmonisation involves the convergence of domestic legislative frameworks rather than their uniformity, and this convergence is strictly dependent upon the national implementing measures (NIMs) [76].

An example of this process can be seen in the transposition of Article 3(2) of Directive 2004/38/EC. The United Kingdom (UK) provided detailed guidance on the transposition of EU directives into national law and employed both the “copyout” (direct adoption) and “elaboration” (adaptation) approaches [77]. In contrast, Spain has not had any comprehensive guidelines, which has led to ambiguity and criticism from legal professionals.

Article 3(2) of Directive 2004/38/EC, which deals with the right to free movement and residence for extended family members of EU nationals, serves as a case study for examining these transposition challenges. A horizontal comparison of the Spanish and English versions reveals sig-

nificant syntactic and terminological similarities due to the “synoptic approach” being taken to ensure uniform interpretation. Every version of a text in different languages maintains the same page count, textual structure, and sentence length, ensuring that identical information appears at corresponding points. Employing punctuation to segment the text into smaller meaning units aids in achieving synchrony, facilitating citations, and enhancing interpretation [78]. However, certain expressions, such as *en el país de procedencia* (the country they come from), exhibit differences that can lead to varied national interpretations. The Spanish version closely aligns with the French one, possibly indicating the use of French as a reference language. Key concepts like *miembro de la familia a cargo* (dependant) and *pareja* (partner) are retained in both versions, but the lack of clear definitions leads to divergent national interpretations. For instance, “dependant” is interpreted financially in the UK, while Spain interprets it as meaning both financial and physical dependence.

The transposition of Article 3(2) in Spain and the UK illustrates these challenges. Spain clarified ambiguous concepts to include both the country of origin and non-origin into relevant definitions, while the UK initially restricted the term to the European Economic Area (EEA) state where the national resides, later amending it after a CJEU ruling. The term *pareja* (partner) was transposed differently: Spain used *pareja de hecho* (unregistered partner), aligning with the Directive’s terminology, while the UK contrasted “partner” with “civil partner”, reflecting a different recognition of partnerships. Initially, Spain only recognised registered partners but later included unregistered partners in its transposition. Overall, the transposition of Article 3(2) underscores the complexities and variations in implementing EU directives within different national legal systems.

A second example of diverging domestic implementing legislation was analysed in the contexts of the CrossJustice [79] and FACILEX [80] projects. The CrossJustice Project aimed to enhance the accessibility and effectiveness of justice in Europe by focusing on key EU criminal justice legislation — including Directive No. 2016/343 on the strengthening of certain as-

pects of the presumption of innocence and of the right to be present at the trial in criminal proceedings — and support practitioners in the identification of the EU procedural rights of persons accused or suspected of a crime. It has provided online tools, databases, and training modules to help legal professionals and citizens understand and apply these directives. Building on this, the FACILEX project advances legal practitioners' access to EU legal standards, emphasising digitalisation and legal technology, covering legislation including the European Framework Decision on the European Arrest Warrant No. 2002/584/JHA. The project leverages AI for the development of tools that support the implementation of EU laws, thus reducing legal uncertainty and enhancing judicial efficiency. Together, these projects are intended to ensure cohesive application of EU law, aiding both legal practitioners and citizens in effectively navigating the justice system. One of the outcomes of these innovative projects is the creation of complete transposition tables that align EU legislation with their national implementing legislation. These tables clearly identify divergences in national legal frameworks, despite the harmonising function that EU legislation should have.

The European Arrest Warrant (EAW) Framework Decision, adopted on 13 June 2002, is a pivotal instrument designed to enhance judicial cooperation in criminal matters within the European Union. Despite its harmonising intent, significant divergences in implementation have emerged across member states. For instance, Article 1(3) of the EAW mandates respect for fundamental rights and fundamental legal principles, a requirement explicitly incorporated into the national legislation of Germany and Italy, thereby providing clear legal safeguards. In contrast, France relies on judicial interpretation to uphold these principles, with the French Supreme Court allowing the refusal of an EAW if it violates fundamental rights. Article 2 specifies that an EAW may be issued for acts punishable by a custodial sentence of at least twelve months or, for sentences already passed, at least four months. This provision is uniformly implemented in the national laws of France, Germany, and Italy, ensuring adherence

to the EU’s custodial sentence thresholds. Article 3 outlines mandatory grounds for non-execution of an EAW, including amnesty, double jeopardy, and age-related criminal responsibility. France, Germany, and Italy have all transposed these grounds into their national laws, though variations in legislative language exist. France and Germany provide detailed statutory provisions, while Italian implementation is comprehensive but may exhibit slight terminological differences. The primary divergence lies in the safeguarding of fundamental rights, with French judicial oversight contrasting with the explicit statutory protections in Germany and Italy. These differences reflect the varying national legal cultures and their impact on the EAW framework’s consistency within the EU’s legal landscape. Understanding these divergences is crucial for assessing the efficacy and uniformity of the EAW across member states.

Last, but not least, EU regulations, though directly applicable and uniformly binding across all member states upon their enforcement, can still result in implementation discrepancies due to varying national administrative practices and judicial interpretations. The General Data Protection Regulation (GDPR) is a pertinent example. Despite its direct applicability, the GDPR has seen various enforcement actions and interpretations by national data protection authorities, which influences the degree of uniformity in its application across the EU [81].

1.2.3 EU law-making and linguistic challenges

Linguistic uncertainty within the EU legislative process is a multifaceted issue, manifesting both intra-lingually, within a single language, and inter-lingually, across multiple languages [82]. This duality is especially pronounced in the EU’s multilingual context. The inherent uncertainty of any language (intra-lingual uncertainty) is compounded by the translation between languages (inter-lingual uncertainty), significantly complicating legal interpretations [83].

The drafting of EU legislation involves multiple languages and is sub-

ject to imperfections due to the presence of non-native drafters, amendments proposed in varying languages, and the inherently negotiated nature of the legislation [56]. These factors can lead to translations that might not ‘sound natural’ or that use phrasing reminiscent of national legislation, which some critics attribute to translator errors. However, translation within the EU involves making choices that often reflect nuanced, non-binary decisions rather than clear-cut errors [53, 84].

Further complicating matters, historical assumptions about the quality of multilingual EU law have changed dramatically post-1992, following significant political integration steps, including the Maastricht Treaty entering into force [85]. The quality of EU legislation has since become a critical concern, leading to initiatives aimed at improving legislative drafting and addressing the EU’s perceived “democratic deficit” [86]. These efforts underscore the intricate balance between ensuring legal clarity and managing the linguistic diversity inherent to EU legislation.

EU translation practices pose unique challenges that traditional translation theories struggle to address due to the dynamic and hybrid nature of EU legal texts. These texts are drafted and translated simultaneously, requiring ongoing adjustments that standard translation processes do not typically accommodate [87]. This complex scenario not only affects the drafting phase but also carries implications for legal interpretation and implementation at the national level, potentially leading to litigation and other political repercussions [88].

The next section of this thesis will explore these challenges in depth, analysing how linguists in the EU navigate these complexities in lawmaking processes.

Linguistic challenges

EU translators must navigate complex linguistic challenges to make legislative texts accessible to diplomats, lawyers, technical specialists, and the general public. The need to balance precision, clarity, and fidelity to the

original text introduces significant linguistic uncertainty [72]. Within the multilingual framework of the EU, where regulations are mostly directly applicable in the national legal systems, the complexity of managing linguistic interactions increases notably [7, 57].

Linguistic indeterminacy, inherent to all languages, is amplified in EU legislation by multilingualism, which introduces inter-lingual uncertainty through grammatical, lexical, and error-induced ambiguities across language versions [89, 82]. Although translation errors are often seen as separate issues, they are intrinsically linked to the broader challenges of multilingual lawmaking, including ambiguity, terminological differences, and textual features like syntax and punctuation [90].

Semantic indeterminacy in legal texts can be either deliberate, to accommodate diverse legal and political environments, or inadvertent. EU legal concepts are intentionally kept vague to ensure applicability across different EU member states, and they reflect political compromises. However, this often results in translations that do not naturally align with the linguistic structures normally used in the member states, leading to divergences in language versions and challenges in maintaining legal consistency [14].

The multilingual EU environment complicates the translation process, demanding an exceptional level of cooperation among translators, lawyer-linguists, and national experts. The creation and translation of EU-specific terminology, a vital component of EU law, requires meticulous coordination to ensure that new legal concepts are understood in a consistent manner across all official languages [91]. This task is compounded by the dynamic nature of EU law, which frequently introduces new terms that must be integrated into the existing legal framework — a condition that can result in lexical indeterminacy [82].

Translators must also deal with systemic differences between languages and legal systems, which can lead to significant terminological discrepancies [92]. The concept of linguistic relativity, among other theories, high-

lights how language informs thought, further complicating the translation of legal terms across diverse legal cultures [93].

In conclusion, the interplay between EU law and national legal systems, coupled with the intricacies of multilingual translation, poses unique challenges [33]. These challenges necessitate a careful approach to term formation and translation to ensure that EU legislation is both effective and comprehensible across different linguistic and legal contexts [94]. The EU is indeed shaping a unique legal culture, anchored in the *acquis communautaire*, which is aimed at building a specialised supranational network of legal concepts. Yet, a challenge emerges as EU legislation is crafted in multiple languages where legal terminologies have established connotations. This framework not only introduces novel terminologies but also incorporates existing national terms, which may overlap in meaning but differ in detail. Consequently, a single legal term could carry multiple interpretations: one specific to a national legal system and another within the EU legal framework [95].

The complexity increases because each application of a legal term in a norm contributes to defining its meaning, and when various legal systems interpret the same term differently, disparate concepts can emerge under the same label [96]. To address this issue, the CJEU has enshrined the duty of consistent interpretation, which consists of “the obligation of national courts and administrative authorities to interpret the applicable national law as much as possible in a way which ensures the fulfilment of obligations deriving from European law” [97]. Nevertheless, the duty of consistent interpretation does not entirely prevent national courts from adopting a different interpretation of legal terms and, consequently, of legal norms, especially in legislative areas where the EU lacks legislative competence. The CJEU indeed refrains from comparing these concepts across EU member states’ legal systems. It holds that, unless expressly stated, terms in EU treaties and legislation should carry an autonomous meaning specific to the EU legal order. Definitions derived from national laws are generally unsuitable if they contradict the aims of EU law or compromise

its uniform application and effectiveness [98].

1.3 The meaning of EU law

Unlike sentences, the semantic value of sentences is not merely determined by its constituent elements and structure; rather, it necessitates more intricate cognitive engagement and context for comprehensive understanding [99]. The discipline of legal interpretation, historically a topic within the philosophy of language, has been increasingly integrated into the philosophy of law, and is particularly influenced by doctrines like consistent interpretation or indirect effect. As previously outlined, consistent interpretation requires domestic courts to interpret national law so that it is consistent with EU law as far as possible [97]. Conversely, the principle of indirect effect describes how EU law can influence, primarily through the entry into force of EU directives, even in the absence of specific implementing measures [100].

The unique interpretative challenges posed by the multilingual nature of EU law have not yet been thoroughly theorised, particularly how they pertain to the broader theoretical implications of meaning and interpretation. This section explores a theoretical approach that integrates legal theory and philosophy of language, specifically semantics and pragmatics.

Legal theory and philosophy of language both play crucial roles, with the former primarily concerned with legal interpretation and the latter with broader textual meanings and interpretations. While philosophers of language have delved into legal interpretation, they have not specifically focused on the context of EU law. The semantic approach to EU law can illuminate the core problems of legal interpretation and influence interpretative choices, thereby not only advancing legal theory but also enhancing our understanding of language comprehension challenges [101].

Several years ago, Timothy Endicott argued that the theory of meaning and interpretation applicable to language in general closely mirrors that

of legal language [102]. Legal language, combining everyday terms with specialist legal terminology, should aim to be accessible to the general populace. In legal settings, judges are tasked with deciding the correct interpretation of legal texts. On the other hand, linguists in academia view multiple interpretations as phenomena to be analysed rather than judged [103].

Despite the rich philosophical discourse on the influence of ordinary language on legal interpretation, the concept of “meaning” in legal contexts is often overlooked [104]. In general, meaning can manifest in several forms — literal, ordinary, linguistic, contextual, communicative, pragmatic, intended, reasonable, and previously interpreted, among others [105]. In a much more contextualised framework, the legal meaning predominantly arises from literal interpretation of the law and the interpretations provided by competent authorities. This belief forms the backbone of the legal informatics methodologies and approaches proposed in the second part of the thesis.

The complexity of meaning is exacerbated in EU law by the presence of numerous authentic language versions of each legal text, which must convey a unified meaning despite linguistic diversity [69]. This challenge underscores the ongoing debate between ‘literalists’ and ‘contextualists’. Literalists believe in a fundamental, abstract literal meaning as the basis for interpretation, whereas contextualists argue that context directly informs meaning [103].

Legal interpretation also involves a distinction between the communicated meaning of a text and its authoritative legal meaning, which can diverge based on judicial interpretation [106]. Thus, while legal texts convey a range of meanings, interpretation often involves pragmatic considerations that transcend the textual content [107]. The dynamic interplay between various types of meaning highlights the complexity of legal interpretation, necessitating a nuanced understanding of both linguistic and legal principles.

A detailed analysis of the philosophy of language dealing with “the meaning of meaning” falls outside the scope of this thesis. Instead, this work primarily identifies and describes legal informatics methodologies designed to support the analysis of legal harmonisation, from a partly multilingual perspective, and legal interpretation, considering the most relevant methods of interpretation employed by the CJEU.

1.3.1 Interpretation of EU law

The Court of Justice of the European Union (CJEU) plays a pivotal role in ensuring the consistent interpretation and application of EU law, as mandated by Article 19 of the Treaty on European Union (TEU). The CJEU is divided into two distinct entities: the Court of Justice and the General Court. The former comprises twenty-seven judges, representing each EU member state, while the latter is staffed by two judges per member state. Supporting the Court of Justice are eleven advocate generals, who contribute to the court’s deliberations by presenting reasoned opinions on the cases assigned to them [108].

The jurisdiction of the CJEU encompasses a variety of judicial proceedings, which are largely determined by the nature of the claims and the identities of the parties involved. References for preliminary rulings can be considered the most particular proceedings. They are sought when domestic courts are uncertain about the interpretation or validity of EU legislation [109]. Pursuing a different objective, infringement proceedings are initiated against member states alleged to have failed in fulfilling their EU law obligations arising from both primary and secondary law. The CJEU also entertains actions for annulment aimed at invalidating EU acts that purportedly contravene EU treaties or fundamental rights, and they can take actions against failure to act, where EU institutions are accused of neglecting their duties. Lastly, actions for damages are pursued when individuals or entities claim to have suffered harm due to the actions or inactions of the EU or its personnel [110].

The CJEU is commonly conceived of as the supreme adjudicator within the EU legal framework. Domestic courts also play a crucial role within this multilayered judicial system by enforcing EU law, which has mostly direct effect on the legal systems of the member states [111, 112]. The interaction between the CJEU and national courts is anchored in the principle of sincere cooperation, which obligates both the European Union and the member states to adopt measures, either general or specific, that ensure the fulfilment of obligations arising from the treaties or from acts of the EU institutions, pursuant to Article 4(3). More specifically, this principle imposes a duty on national courts to interpret domestic law in a manner that is aligned with EU law, even when such an interpretation reveals discrepancies between national and EU norms [113].

Despite the overarching requirement for consistent interpretation, national courts retain the discretion to deviate from interpretations rendered by the CJEU. Rather, the principle of sincere cooperation envisages a scenario where national courts and administrative authorities are obliged to interpret applicable national laws in a way that facilitates the fulfilment of obligations under EU law. This necessitates a comparative legal approach that not only considers how EU norms are applied by the CJEU but also how they are interpreted at the national level.

As already mentioned, salient examples of the interpretative challenges posed within the EU framework is the definition of “spouse” and “family” under the EU Citizens’ Directive. The interpretation of EU law is pivotal in shaping the legal landscape across its member states, ensuring uniformity and consistency in the application of laws. This interpretative process is governed by a set of established principles and is faced with unique challenges stemming from the EU’s multilingual framework. According to legal doctrine, there are four classical principles underpinning legal interpretation by the CJEU, namely textual/literal interpretation, contextual/systematic interpretation, teleological interpretation and comparative interpretation. [23]. The incorporation of the comparative method adheres to the classification in [114], while historical discussions

on the diachronic evolution of legal norms are subsumed under the broader systematic category. Other authors have proposed different classifications, which are outside the scope of this thesis. [89, 33].

Literal interpretation

As an interpretative principle that crosses multiple legal systems, the CJEU relies on the literal interpretation method, a foundational criterion based purely on the semantic and syntactic content of the wording of the law, including the norms being interpreted [109, 115]. Literal interpretation can be considered a pillar of the principle of legal certainty, since it fosters the predictability of the CJEU's judgments by looking at the common meaning of the words used in the norm [7]. However, according to the maxim *interpretatio cessat in claris*, the literal approach loses its usefulness where the norm is not clear and precise, especially considering that the treaties are imbued with a “purpose-driven functionalism” that is suitable to teleological interpretation [116]. Within the EU, the issue of clarity and precision assumes an outstanding importance. As the Union is composed of twenty-seven member states, a norm can be fully clear in one linguistic version and not so in another one. [109].

Systematic or contextual interpretation

The second method is systematic or contextual interpretation, which refers to the context of the norm i.e. its historical background, the legal source where it was codified, how it fits within the law, the procedure that resulted in the enactment of the specific EU law provision in question etc. [98]. It is based on the premise that the legislator is a rational actor, meaning that the authors of the treaties have established a consistent and complete legal order in which each norm has a specific *effet utile* [117]. As the historical background of a law is a relevant aspect of the systematic approach, the CJEU over the last decade inverted the negative trend [118] in using the so-called *travaux préparatoires* as a source of interpretation. Notably, the

CJEU made increasing use of the preparatory works of the EURATOM, Maastricht, and Constitutional Treaties, which have played a role in recent judgments and have become a source of EU law [119].

Furthermore, the preambles of laws — composed of recitals that illustrate the path followed by the legislative measure under scrutiny, the policy initiatives underpinning its adoption, and much more information — have been constantly getting bigger over the years and are steadily used by the CJEU for hermeneutic purposes [120]. In the preambles of European legislation, provisions that have not been allocated legally binding status within the articles of the respective law are frequently included as recitals. This arises from failure to achieve common political consensus during the legislative procedure, often among representatives of the EU member states [120].

Teleological interpretation

More generally, where a EU law provision is open to several interpretations, preference must be given to the one that fosters its effectiveness the most. In this respect, the need to focus on the *telos* (a Greek term meaning aim, objective) of the norm is fundamental and shows that the teleological and systematic approaches are often interlinked, since the latter can be useful for understanding the real purpose of a provision [121]. The CJEU consistently resorts to analysing the goals of a provision in order to interpret its content effectively; the objectives pursued by a norm can be contained in the law itself or also in the founding treaties of the Union [122].

Teleological interpretation can be considered the distinctive interpretative method of the CJEU, since the treaties are imbued with a purpose-driven functionalism. This means that the founding treaties are based on the idea that there are objectives of paramount constitutional relevance that the EU must pursue and implement [122]. Bengoetxea [123] identified three types of teleological interpretation, i.e. functional interpretation, aimed at pursuing the *effet utile* of the provision, teleological interpreta-

tion *stricto sensu*, which builds on the objectives pursued by the norm in the event that it has an ambiguous meaning, and consequentialist interpretation, which assesses the possible consequences of an interpretative choice [124]. Moreover, given the fact that a provision ordinarily pursues more than one objective, the CJEU applies the principle of proportionality or reasonableness when assessing the prevailing purpose in the specific case at stake.

Comparative interpretation

Lastly, although empirically less significant, comparative interpretation is a consequence of the multilingual nature of the EU, where legislation is drafted in twenty-four official languages and the principle of equal authenticity in the EU requires that all language versions of legal texts are treated as identical in meaning despite their linguistic differences [7]. This presents significant challenges in legal interpretation, because no single language version can singularly convey the law's meaning. It must be derived collectively from all versions [125]. Despite their equal status, the inherent variability in natural language and translation approximations means that different language versions often diverge in meaning [126].

This divergence is exacerbated by the collective and multicultural nature of EU legal drafting, often involving compromises that prioritise consensus over clarity [127]. These factors contribute to inevitable discrepancies among language versions, complicating traditional legal interpretation methods and raising questions about the legislator's intent [128]. When differences or incompatibilities between language versions arise, determining the authoritative interpretation becomes particularly complex.

In summary, the interpretation of EU law by the CJEU is a sophisticated process that balances textual, contextual and comparative methodologies to achieve coherent and effective application of laws across the Union. This process not only promotes legal uniformity but also adapts EU law to the changing needs and complexities of its member states. The

principles of interpretation employed by the CJEU thus play a crucial role in the integration and functionality of the European legal system, highlighting the dynamic interplay between the law and the EU’s policy objectives. As a last principle underpinning EU legal interpretation, it is essential to point out that while not strictly bound by the doctrine of *stare decisis*, which “involves a court’s choice to stand by a precedent notwithstanding suspicions (or worse) about its wrongness” [129], the CJEU promotes legal consistency and certainty by referencing its prior rulings. This practice helps maintain coherent legal interpretation across various cases, enhancing predictability for EU citizens and member states. In this regard, the approaches to legal informatics presented in the second part of the thesis emphasise the need to incorporate rulings and interpretations from the Court of Justice of the European Union (CJEU) into the realm of legal knowledge under examination.

Following this perspective, the next chapters will address the contribution that legal informatics methodologies and tools can provide to support legal interpretation, considering the difficulties that not only judges but also legal consultants, lawyers and analysts experience when they need to manage a vast amount of legal sources or when tasked with identifying the precise meaning that concepts, judicial citations and arguments assume once applied to the concrete cases under scrutiny.

1.4 Digital tools for EU legal analysis

The Publications Office of the EU provides an Inter-Institutional Style Guide for each language, which is crucial for maintaining consistency and historical continuity in EU translations. The style guides are particularly valuable when integrating new languages into EU documentation as they form new patterns and terminology, often based on English or (in the past) French texts [130].

Translators in the EU utilise a variety of tools to support their daily

work, categorised into terminology databases, translation memory technology, and machine translation. Notably, the InterActive Terminology for Europe (IATE) serves as a comprehensive terminological database, drawing contributions from all European institutions and accessible to the public. It contains over 8 million terms and is pivotal for ensuring high-quality, consistent translations across EU languages [131]. Another critical tool is Euramis, which utilises a central translation memory to streamline the translation process and ensure coherence across documents. However, it requires careful management to maintain consistency, particularly given its broad usage across various types of documents [7].

Especially important for the automatic analysis of EU legal documents is the Cellar database [132], managed by the Publications Office of the European Union, which provides access to EU publications and their metadata through RESTful (representational state transfer architectural style) web services and a SPARQL (SPARQL Protocol and RDF Query Language) interface [133]. It organises data using the Common Data Model (CDM) based on the FRBR (Functional Requirements for Bibliographic Records) hierarchy (Work, Expression, Manifestation, Item). The database supports multiple views, uses standardised models and vocabularies, and ensures unique identification of publications. APIs (application programming interfaces) for RSS (Really Simple Syndication) and Atom feeds allow users to receive updates on new or existing publications. This repository is essential for scholars interested in collecting vast amounts of EU legal documents for automatic analysis and classification. The aim of providing access to European Union documents in structured formats like XML (eXtensible Markup Language) and JSON (JavaScript Object Notation) is to achieve semantic interoperability, a priority for the European Commission. It was recently codified in EU Regulation No. 2024/903, the so-called Interoperable Europe Act, with its binding requirements on domestic public administrations concerning, among others, access to documents and data standards. The main goal of semantic interoperability is to facilitate communication between national public administrations without

losing semantic content. To that end, the EU Commission has proposed adopting the Akoma Ntoso standard in a version tailored for European law [134]. The EU Commission has already adopted other more widespread standards, including XML and JSON.

Additional resources include ELISE (European Institutions Linguistic Information Storage and Exchange), a documentation system that tracks the language issues encountered and resources used across different stages of the legislative procedure, enhancing collaboration among institutions [72, 135]. EUR-Lex provides another essential resource, offering access to all EU legislation and case law. These documents are increasingly integrated into translation memories to ensure better phraseological and terminological consistency [72, 135].

The EU’s commitment to developing sophisticated linguistic tools like Quest — a metasearch tool combining multiple databases for efficient terminology searches — highlights ongoing efforts to optimise translation practices within the multilingual framework of the EU, ensuring clarity, consistency, and legal integrity across its 24 official languages [136].

In late 2017, the European Commission launched a new translation service. eTranslation is a cutting-edge neural machine translation service provided by the European Commission’s Directorate-General for Translation (DGT) to replace the earlier MT@EC system [137]. This service leverages advanced AI tools, specifically neural machine translation (NMT) techniques, which employ artificial neural networks to predict the likelihood of sequences of words. Neural networks surpass older statistical methods in predicting word sequences because they can learn complex patterns from large data sets, handle context through architectures like RNNs, and use word embeddings to understand semantic relationships. Their ability to learn end-to-end and scale with data improves performance and adaptability across various NLP tasks, making them more effective at generalizing from seen to unseen data. This results in translations that are more accurate and natural-sounding compared to older statistical methods. eTranslation supports translations to and from any of the 24 official EU languages, as

well as Arabic, Chinese, Icelandic, Japanese, Norwegian, Russian, Turkish, and Ukrainian. It is designed for use by EU institutions, public administrations, universities, small and medium-sized enterprises (SMEs), and other eligible entities across the EU, Iceland, Norway, Liechtenstein, and Ukraine. With high security protocols ensuring that all data remains within the Commission's infrastructure, eTranslation offers translations in various domains, such as EU legal language, general text, and finance, and supports all common Microsoft Office formats and their equivalents. Although the raw machine translations offered are intended for quick understanding or as a starting point, professional translators are encouraged to revise these translations to ensure high-quality outputs.

Building on the foundational principles of EU legal interpretation, the traditional methods of legal informatics — legal ontologies, similarity metrics, network analysis, harmonised glossaries etc. — are highly effective tools that align well with the interpretative techniques employed by the CJEU. Ultimately, these methods are useful for analysing the phenomenon of legal harmonisation in the European Union, which is closely intertwined with judicial interpretative methods.

Lastly, harmonised glossaries and indices of harmonisation serve as valuable assets. When integrated into a unified tool, they facilitate the simultaneous application of multiple interpretative methods, bridging gaps between different linguistic and legal interpretations. This integration promotes a more holistic approach to legal interpretation, ensuring that the multiplicity of EU legal texts are interpreted in a way that is both comprehensive and coherent. Such technological integration in legal informatics not only aligns with but also enhances the interpretative strategies of the European Court of Justice, fostering greater precision and uniformity in the interpretation of European law.

Chapter 2

Similarity metrics for literal and comparative interpretation

Similarity metrics are fundamental to the comparative interpretation of legislative texts, enabling precise analysis across different linguistic versions, but also intra-lingual comparisons. Moreover, these metrics are crucial in literal interpretation, enhancing the accuracy and clarity of legal statements. Although the impact on literal interpretation might not be immediately apparent, further explanation reveals how this approach contributes to a deeper and more accurate understanding of legislative language. The objective is to ensure consistency and coherence across multiple languages. By employing similarity metrics, legal analysts can perform direct, text-based analysis that helps to identify and reconcile linguistic discrepancies. This not only enhances the clarity and uniformity of EU laws but also contributes to the overall effectiveness of legal translations, promoting better understanding and implementation across diverse legal systems. Besides their general effectiveness in enabling the comparison of text portions, similarity metrics are foundational to most machine learning approaches for legal information retrieval, analytics, and text generation. These metrics enhance the ability to automatically analyse and

generate legal texts, ensuring consistency and accuracy across various applications [138].

2.1 Related work

Recent advancements in legal research reveals a burgeoning interest in deploying artificial intelligence (AI) methodologies for enhanced legal analysis [139, 140, 141, 142, 143, 144, 145, 146]. Computational text analysis has emerged as a critical area within legal research, with most recent applications extending across diverse sub-areas such as case-law analysis [138, 147], legal process management [148], and legal text generation [149]. Applications in this area typically target sources of unstructured text including online reviews, social media posts, and various other forms of digital commentary [150, 151], with the aim of revealing patterns and deriving insights via both supervised and unsupervised learning techniques [152].

In the legal sphere, the adoption of Natural Language Processing (NLP) has been particularly rapid, with Legal Informatics emerging as a vital sub-field focused on the application of Information and Communication Technologies (ICT) to legal contexts [153, 154]. Legal sources present unique challenges and opportunities for computational analysis [155] because the text often follows a formal structure and includes elements such as preambles and citations.

Within AI, machine learning techniques have come to include similarity measures as a pivotal component of NLP workflows [156]. Methodologies for identifying similarities in legal documents typically fall into two categories: (i) text-based techniques that analyze the contents of the documents themselves [157]; and (ii) network-based approaches that draw on citation networks [158]. This chapter focuses on the former category, particularly recent studies concerned with the alignment between EU directives and their national transpositions [159, 160], while the third chapter addresses network-based approaches.

NLP and text mining methodologies can be instrumental in helping the European Commission and legal practitioners to analyse the transposition of directives in detail [161]. Recent initiatives have utilized embedding models [162] to represent legal texts within semantic vector spaces, and Cosine Similarity (CS) measures [163] to compare these representations. Innovative research has also explored the use of graph-based methods that pinpoint key concepts and representative sentences for the purpose of determining thematic similarities in court rulings [164]. Experiments on legal judgments have shown the effectiveness of the Term Frequency - Inverse Document Frequency (TF-IDF) approach, particularly when focusing on legal-specific vocabulary [165]. Previous works have demonstrated the significance of combining TF-IDF, stemming, and co-occurrence networks for automated analysis of legal texts [166]. In summary, recent analysis has revealed that traditional methods like TF-IDF and Latent Dirichlet Allocation (LDA) tend to outperform more recent context-aware approaches such as Bidirectional Encoder Representations from Transformers (BERT) and Law2Vec in computing document-level similarities [167].

2.2 Comparing European procedural safeguards in criminal proceedings: a case-study

The safeguarding of fundamental rights within the European Union (EU) is significantly influenced by the diverse legal frameworks of its member states which affect, among other matters, the effectiveness of procedural rights [168, 169]. Legal practitioners face challenges in determining the applicable legislation and corresponding procedural rights for individuals accused or suspected of criminal activities, and these complexities are exacerbated by linguistic diversity and the unique characteristics of each national legal system [170, 171]. In 2009, the EU Stockholm Programme, a strategic plan for 2010-2014, placed significant emphasis on strengthening judicial

cooperation across the European Union. One of its primary objectives was to ensure that legal decisions and protections were recognised and enforced uniformly throughout all EU member states. This included enhancing mutual trust between national legal systems, facilitating cross-border legal procedures, and ensuring that citizens could rely on consistent judicial protections regardless of where they were in the EU. The programme sought to create a more interconnected and efficient judicial area, where cooperation in civil and criminal matters would be seamless, thus promoting a secure and just Europe. Even after the adoption of the Stockholm Programme and related EU Directives on procedural safeguards for accused persons in criminal proceedings [172], these challenges persist [173]. Indeed, European directives enter into force only after being transposed into the national laws of the member states via NIMs [174], which are domestic legislative acts. There are two predominant approaches to transposing EU law into national legislation: i. *copy-out*, technique commonly used in the UK [175] where the implementing legislation either replicates or closely follows the original text of the directive, and ii. *elaboration*, where legislators interpret the directive’s provisions in such a way that they can be integrated more clearly into national law. The *copy-out* strategy is the most commonly used, so it is likely that the texts of the NIMs will closely resemble the original directive. Accordingly, textual similarity is expected between NIMs and between NIMs and EU legislation. This chapter examines *approximation of laws* and *legislative harmonisation*. Both these terms refer to alignment between EU and national legislative frameworks. Legislative harmonisation within the EU serves a dual purpose. First, it reduces legal disparities across member states to enhance economic, social, and cultural exchanges. Secondly, it seeks to achieve various political objectives, such as establishing a European single market, setting common minimum standards for social protection, and regulating the rights of suspects and accused persons in criminal proceedings.

In summary, the primary aim of this chapter is to evaluate the effectiveness of conventional text analysis methods on NIMs. The specific research

questions are:

- How effectively can the implementations of EU directives across various States be compared using NLP techniques?
- When analysing the “explicitly transposed” articles of each directive, is it feasible to utilise certain analytical metrics, such as similarity or network measures, to compare the NIMs? Additionally, are these metrics effective when applied at the granularity of individual articles?

2.2.1 The CrossJustice project

This thesis draws upon the CrossJustice (CJ) project¹, which examined whether national instruments implementing EU directives were aligned with the *acquis communautaire*, specifically in safeguarding the fundamental rights of individuals accused or suspected of a crime — a key objective of EU justice policy. The project engaged legal experts in evaluating the congruence between national legal frameworks following the enactment of six EU directives. A key output of this project was the development of a web platform designed to facilitate the dissemination and application of the findings.

The CJ project addressed the aforementioned challenges by leveraging ICT to pinpoint crucial discrepancies and proposing solutions from a comparative standpoint, enhancing the functionality and collaborative potential of judicial systems. The associated online platform offers comprehensive support on procedural rights. It is freely available and is intended for legal professionals primarily, but also accessible to law students, non-governmental organisations (NGOs), and indeed all EU citizens. The CJ platform², lastly updated in 2022, offers: i) up-to-date, complementary information and advisory services tailored for legal practitioners including

¹<https://www.crossjustice.eu>

²<https://www.crossjustice.eu/en/index.html#crossjustice-platform>

lawyers, public prosecutors, and public administration employees, but also available to law students and general citizens; ii) capacity-building opportunities, including a broad variety of tools and publications, tailored for legal professionals and law students.

2.2.2 Manual annotation of directives

National implementations were classed with four alternative distinct labels by legal experts, each label representing a different type of implementation strategy. These labels were systematically applied to classify the approaches each member state employed to incorporate EU directives into their national laws. This structured annotation scheme helps to provide comprehensive analysis of the transposition methods used across different jurisdictions. The labels are as follows:

1. **explicit transposition:** This approach to integrating EU directives involves adopting new laws or modifying existing ones.
2. **de facto/indirect implementation:** In cases where the rights or provisions stipulated by the directive already exist within the national legal framework, explicit transposition may be deemed unnecessary.
3. **absence of national implementation:** This happens either because there is no national legislation explicitly or indirectly transposing the directive, or because existing laws fail to meet the requirements of the directive.
4. **non-requirement of specific transposition:** Transposition may not be required if: i) the legal provision in question has no deontic or constitutive significance. Examples are articles that merely delineate the directive's scope; ii) there is a possibility for member states to opt-out of specific provisions. One example is Article 6(3) of Directive 2016/800.

Table 2.1: Distribution of NIMs segments for the specified EU directives across member states, categorised into four types: explicitly transposed (Explicit), de facto/indirectly implemented (Indirect), no national implementation (NoImpl), and specific transposition not required (NotReq).

member state	Explicit	Indirect	NoImpl	NotReq	Total
<i>Bulgaria</i>	40	151	17	15	223
<i>Croatia</i>	146	81	24	0	251
<i>France</i>	49	153	41	0	243
<i>Germany</i>	99	234	11	0	344
<i>Italy</i>	65	221	32	0	318
<i>Netherlands</i>	150	146	57	16	369
<i>Poland</i>	32	154	86	0	272
<i>Portugal</i>	0	353	22	0	375
<i>Romania</i>	85	239	50	0	374
<i>Spain</i>	91	135	89	0	315
<i>Sweden</i>	8	325	0	41	374
Total	765	2,192	429	72	3,458

2.2.3 Overview of the legal corpus

The six EU directives examined (No. 2010/64, 2012/13, 2013/48, 2016/343, 2016/800, 2016/1919) are implemented differently across various EU member states. Legal experts participating in the CJ project annotated segments of each directive — articles or paragraphs — with appropriate labels as well as the corresponding texts of the national transposing legislative provisions along with commentary from the TT under scrutiny published on the CrossJustice platform. On June 1st, 2021, the TT contained 3,458 annotations showing variations in implementation among the member states. For instance, Bulgarian legal experts registered the fewest annotations (223), while Portugal ones registered the most (375), underlining the complexity of transposing EU laws into national legal systems. The extent of explicit transpositions also varied, with the CJ table showing Croatia and the Netherlands having the highest level of direct transpositions, and Portugal and Sweden the lowest.

2.3 Methodology

The methodology employed for this research had two distinct stages, each critical to the analysis. The first stage involved comprehensive text processing. This stage was dedicated to preparing the textual data for analysis, which included cleaning and normalising the text to ensure consistency and removing any irrelevant or extraneous information. Techniques such as tokenisation, removal of stop words, and stemming are commonly applied during this stage to refine the data further, and was carried out for this research.

Following the initial preparation, the second stage focused on measuring text similarity. This involved applying various computational techniques to quantify the degree of similarity between texts. Methods such as Cosine Similarity, the Jaccard Index, and other statistical measures were utilised to evaluate and compare the textual data. This step is crucial for identifying patterns, relationships, and differences across the dataset, thereby providing insights into the underlying structures of the different texts.

2.3.1 Legal text processing

This study adopts a comprehensive NLP pipeline that is composed of several processing steps with the aim to prepare and explore textual data effectively. The initial step involved preprocessing techniques for converting the text into a normalised form. This included converting all characters to lowercase, removing stop words and punctuation, and applying part-of-speech (POS) tagging in order to focus on nouns, verbs, and adjectives.

Following the preprocessing, stemming was employed to reduce words to their root forms using the Porter stemming algorithm [176]. Stems that were merely a single character in length were discarded. The stemming process helps to reduce text variability and enhance the homogeneity of the linguistic data.

The core of the methodology was the modelling of textual data. Texts were represented numerically in vectors constructed from term frequencies. The approach utilised both bag-of-words and n-gram strategies to capture contextual information. Specifically, sequences of up to three contiguous terms (trigrams) were considered in order to obtain a more nuanced representation of the textual data. The *bag-of-ngrams* technique is particularly effective for capturing richer linguistic patterns compared to single-word models.

Additionally, Term Frequency-Inverse Document Frequency (TF-IDF) transformation was applied to weigh the terms within the textual corpus. This widely utilised NLP technique enhances the significance of distinctive terms by emphasising words that are frequent in a document but rare across the document corpus [176]. The TF-IDF metric adjusts the raw frequency of a term (TF) with the term's inverse document frequency (IDF), which is derived from the logarithm of the ratio of the total number of documents to the number of documents containing the term.

To manage the high-dimensional data resulting from text vectorisation, a Document-Term Matrix (DTM) was constructed where each row represented an individual legal document (a NIM) and each column corresponded to a term or stem from the n-gram model. Given the extensive number of features typically generated in such analyses, dimensionality reduction techniques were employed to improve computational efficiency and enhance the interpretability of the data. Multidimensional Scaling (MDS) was implemented for this purpose as it effectively reduces feature space while preserving the distances among data points, which is critical for subsequent analyses of document similarities [177].

2.3.2 Implementation of the similarity metric

In the pursuance of the research objectives, Cosine Similarity (CS) was leveraged to explore text similarity, which is a robust metric commonly employed in this area of research. Mathematically, CS quantifies the co-

sine of the angle between two vectors within a multidimensional space, offering a measure of their orientation similarity. Specifically, the similarity between vectors from two NIMs, labelled A and B, is calculated using the formula:

$$CS(A, B) = \frac{V1 \cdot V2}{\|V1\| \|V2\|}$$

Here, the numerator represents the dot product of vectors V1 and V2, which correspond to NIMs A and B respectively. The denominator, the product of the Euclidean norms of V1 and V2, normalises this measure, ensuring that the similarity value ranges between -1 and 1. The CS metric facilitates comparison of the textual alignment between legal documents, such as the explicit transpositions of Article 1 by different member states. Using this approach, the comparative similarity of each directive and their NIMs was determined, drawing on computations performed by the *Cosine_Similarity* method from the *scikit-learn* Python library's *sklearn.metrics.pairwise* module. The outcome of this analysis was the identification of the most similar NIMs to the directives under consideration.

2.4 Results

This section presents the empirical findings of the investigation into the representation and analysis of legal texts within the framework of the EU directives. Initially, the transformation of legal texts, from explicit transpositions into quantifiable vectors that capture their inherent linguistic features, was explored. Subsequently, the similarity of these representations across different EU member states was assessed in order to understand the degree of harmonisation achieved through transposition. The methodologies employed for text representation included bag-of-ngrams output, dimensionality reduction techniques, and the use of similarity metrics to compare legal documents. These techniques were crucial for interpreting

ID	EUdir	State	NumArt	Label	Text
2039	0013	Croatia	art_2	Exp	summon suspect must specifi suspect suspect theins...
2040	0013	Croatia	art_2	Exp	upon arrest arrest person must immedi provid writt...
2041	0013	Croatia	art_2	Exp	letter right must deliv accus person search warran...
2043	0013	Spain	art_2	Exp	ani person punish act attribut may exercis right d...
2044	0013	Spain	art_2	Exp	admiss complaint suit ani procedur action imput cr...
2045	0013	Spain	art_2	Exp	right defens shall exercis without limit expressli...
2050	0013	Croatia	art_2	Exp	prior file indict compet court bodi proceed perpet...

Figure 2.1: Overview of explicit transpositions (ET) for each portion of Article 2 of European Directive No. 2012/13.

the complex landscape of EU legal texts and they provided a foundation for the visualisations and deeper analyses that followed.

2.4.1 Representation of legal texts

This sub-section describes the process of transforming texts, specifically those texts categorised as *explicit transpositions*, into a fixed-length vector format. The procedure involved utilising a *bag-of-ngrams* approach and implementing dimensionality reduction techniques, as described in the following paragraphs.

Explicit transpositions. The annotations undertaken by legal experts from the CJ project were examined. They marked the segments, including articles and paragraphs, of relevant EU directives that had been explicitly transposed into domestic legislation. For example, EU Directive No. 2012/13 had 563 different implementations; of these, 245 were explicit transpositions. Notably, Article 2 had only 7 such explicit transpositions in two member states — Croatia (4) and Spain (3), as illustrated in Figure 2.1. Further, the merging process of all implementations pertaining to the same part of the directive for each member state was considered.

Bag-of-words and n-grams. All the explicitly transposed texts of relevant NIMs formed the corpus for each EU directive. The textual content of each document was represented numerically using the bag-of-words model,

Table 2.2: Cosine Similarity of implementations by three pairs of member states of Annex 1 of EU Directive No. 2012/13

Member States		Vect100	MDS100	MDS200
France	Romania	0.376	0.055	0.056
France	Spain	0.581	0.206	0.205
Romania	Spain	0.460	0.079	0.084

which builds on a lexicon of all the unique *stems*. Moreover, a more refined *bag-of-ngrams* approach was applied, which grouped *stems* into sequences of varying lengths (*n-grams*). For the national implementations under scrutiny, we processed the *stems* for 1,714 distinct segments. Specifically, the median number of *stems* calculated was 105, with the highest number recorded being 1,365 *stems*. This maximum occurred in the context of Article 10(3) Directive No.2013/48.

Dimensionality reduction. The finalised corpus includes vectors for each article explicitly transposed in the TT, where each ‘column’ represents a term or an n-gram. For example, EU Directive No. 2012/13 is represented by 43 vectors, while the individual stems total 957. In scenarios involving bigrams, the number of features increases to 4,549, and for trigrams, it reaches 6,213. Dimensionality is effectively reduced, using multidimensional scaling, to sizes such as 100 or 200 features.

Similarity. The Cosine Similarity metric is employed to quantify the degree of similarity between the vector representations of texts from two implementations in different member states. A case in point is Annex 1 of EU Directive No. 2012/13, explicitly implemented by France, Spain, and Romania. Manual inspection of the corresponding NIMs indicated a remarkable similarity between the NIMs of France and Spain, and this was corroborated by CS measurements using both the top 100 most frequent terms and the MDS method with 100 or 200 features, as shown in Table 2.2.




Figure 2.2: Heat map visualisation of comparative text similarity between the NIMs of member states, calculated with CS similarity. Lighter shades indicate greater similarity.

2.4.2 Heat map visualisation

This section introduces heat maps as a visualisation tool to enhance comprehension of the similarity metrics calculated between pairs of member states for various parts of NIMs. In these heat maps, darker shades such as blue or green signify minimal similarity, whereas lighter shades like white or yellow denote higher degrees of textual similarity.

For illustrative purposes, Figure 2.2 presents a heat map for NIMs pertaining to EU Directive No. 2012/13. The visualisation distinctly highlights the textual similarities, with lighter shades indicating greater similarity, such as that between France and Sweden compared to other EU member states. The diagonal elements in grey indicate absence of comparison since text from the same State is not compared with itself.

These visualisations are indispensable for swiftly conveying the nuances of textual similarity across NIMs. Feedback from legal experts involved in the project highlighted the utility of heat maps in facilitating analysis of the degree of harmonisation across national legislation, which affirmed their effectiveness as a diagnostic tool.

2.5 Conclusions

This chapter discussed the first outcome of an ongoing research project involving the disciplines of NLP and Law, focusing on the key concept of harmonisation in European law. The main topic under scrutiny is the deployment of computational text similarity techniques to enhance the coherence of legislative norms. This investigation focused on how these norms are adapted within the domestic frameworks of EU member states. The research effort involved similarity metrics, analysis of the results, and a visualisation tool to demonstrate how an established NLP pipeline for preprocessing text and similarity metrics could be applied to support legal harmonisation in different legislative areas.

The next chapter outlines the application of network analysis, based on the co-occurrence of terms or *stems*, with the aim of identifying interrelationships between legislative articles and recitals. The approach has already provided meaningful results [166, 178] in modelling interrelationships between normative provisions [179]. Different similarity techniques, including the Best Match 24 (BM25) term-based ranking model will be explored and hybrid approaches employing graph embedding methods such as Node2Vec or text embedding methods such as Word2Vec, Global Vectors for Word Representation (GloVe), and BERT will also represent future research endeavours. Finally, the evaluation will be extended according to a legal expert assessment of the results and propose an extension to the technology used in the CJ project. Considering the rise of Large Language Models (LLMs) in recent years, one of the most promising research avenues involves developing a methodology based on LLMs and Retrieval Augmented Generation (RAG). This methodological framework aims to automatically create transposition tables using the intrinsic knowledge of LLMs, which is further enriched through an external knowledge base of EU and corresponding NIMs. This approach could help reduce hallucinations and effectively address questions regarding the alignment of domestic and EU legislation.

Chapter 3

Network analysis for contextual interpretation and judicial precedents

After reviewing the main scientific contributions on network analysis, particularly for the legal domain, the following chapter examines three research projects resulting from collaborations between legal experts and data scientists. Network analysis methodologies and legal ontology engineering as a subcategory, are instrumental in systematically organising concepts, definitions, and legal requirements in the context in which they are framed. These tools help to elucidate relational structures and dependencies among legal norms, thereby supporting their application in real-world legal contexts. In the framework of legal harmonisation analysis, network analysis and legal ontologies play a crucial role in modelling instances where European legislation is enforced in member states through jurisprudential mechanisms such as rulings from supreme or constitutional courts. Network analysis offers a variety of analytical metrics that help identify node weights and centralities as well as the most significant relationships within the legal framework. This makes it an effective tool to be used with advanced ML and DL models, providing a robust semantic foundation from which large language models (LLMs) can derive accurate

knowledge. As a result, these models are better equipped to generate more accurate inferences and produce textual outputs that align closely with the specific requirements of user queries. This process enhances the precision and coherence of automated legal analysis, ensuring that the models can effectively address the complexities of legal harmonisation across jurisdictions.

The process of interpreting legal documents, also known as legal hermeneutics, plays a crucial role in adapting legal norms to new and unforeseen circumstances [180]. These documents typically exhibit a degree of vagueness that intentionally allows for diverse interpretations, which adds complexity to legal analysis [181]. Consequently, setting up a robust manual annotation framework is essential for effective machine learning research. Despite its importance, manual annotation remains a challenging task that requires considerable time and effort [182]. Recent discussions have highlighted issues related to the clarity and adequacy of annotation schemes, emphasising the need for well-defined protocols in legal AI research [179].

Further, the application of logical frameworks to model legal reasoning and argumentation has been explored extensively [183, 184]. Prior research delved into the use of ontologies in legal document management systems to enhance the retrieval and interpretation of legal information [185]. A burgeoning field of study is legal text analysis from integrating machine learning techniques, natural language processing (NLP), and network analysis [186]. In particular, information extraction from legal text often focuses on analysing citations through semantics-based network methodologies, as demonstrated in the citation network research of [187].

3.1 Related work

This section delves into three key research areas that contribute to the understanding of network analysis in legal informatics. The first subsection examines methods for identifying and linking norms within legislative

texts, highlighting advancements in text similarity techniques. The second subsection focuses on the application of NLP and ML in legal text analysis, illustrating their effectiveness in various tasks. Finally, the third subsection explores the analysis of legal citation networks, detailing methodologies for extracting and ranking citations within legal sources.

3.1.1 Identifying and linking norms

Legal informatics has made significant strides in identifying connections between normative texts using advanced text similarity techniques. Humphreys et al. [160] explored automated identification techniques to discern implicit relationships between conceptually similar norms, with a particular focus on linking recitals to (sub-)articles in EU legislation. Nanda et al. [188] assessed the similarity between various elements of EU legislation and their NIMs, and evaluated a comprehensive range of similarity measures for this purpose.

Further development in this field is illustrated by the classification framework in [179], which identifies eight distinct types of implicit norm links based on an in-depth examination of Directive 2004/23/EC of the European Parliament and of the Council of 31 March 2004 on setting standards of quality and safety for the donation, procurement, testing, processing, preservation, storage and distribution of human tissues and cells. Complementary investigations have adopted a network analysis approach to explore the complex interplay between norms within systems of interrelated authoritative legislative texts. Sadeghian et al. [189] focused on automated labelling of citations within a legal citation network. Koniaris, Anagnostopoulos and Vassiliou [190] explored the dynamics and multi-layered structures of European legislation networks, emphasising their temporal evolution and complex multi-scale characteristics.

The narrative of this chapter extends to how European directives and regulations — while prescriptive — provide a framework that allows national entities considerable discretion in how they define detailed norms

and procedures that align with the overarching goals of European legislation. This characteristic highlights the principle-oriented approach of directives, which typically favour balanced over defeasible reasoning.

To illustrate the importance for legal interpretation of identifying relationships between different parts of legislative texts, consider the role of recitals in EU legislation, which, illustrate the reasons underpinning the adoption of an act, thus acting as an interpretative legal benchmark [191]. The CJEU has developed this function through case law. As stated in [120], “evidence, albeit indirect, of the importance of recitals in transposition can be found in every case in which the CJEU strikes down a local provision in a transposed rule due to the influence of a recital upon the scope of the transposed legally-binding provision.” Recitals are also used to interpret cases brought by individual claimants. Recitals are allowed to expand, but not restrict, the scope of an ambiguous provision. As mentioned in Chapter 1, the CJEU employs three main methods of interpretation [192]: literal interpretation, which is focused on the wording of the law; systematic interpretation, which considers the context of the norm, its historical background, legal source, and placement within the law; and teleological interpretation, which considers the most suitable way to realise the norm’s purpose.

An example of the use of recitals in systematic interpretation can be seen in cases C-162/97 [193] and C-344/04 [194], where both recitals 14 and 15 of Regulation 261/2004 regarding compensation to air passengers) were used to define the term “extraordinary circumstances” in Article 5 of the Regulation. Article 5 states that the operating air carrier shall not be obliged to pay compensation if it can prove that the cancellation was caused by extraordinary circumstances that could not have been avoided even if all reasonable measures had been taken. Examples of extraordinary circumstances include political instability, meteorological conditions, security risks, and strikes. These examples were used as analogies by the court to help determine the extent to which air carriers are exempt from paying compensation.

The complex interrelationship between recitals and articles remains a contentious issue. Various doctrinal viewpoints exist [120] concerning the interplay between them:

- recitals exert no influence;
- recitals take precedence over normative provisions;
- recitals and normative sections are of equivalent status;
- recitals are subordinate to normative provisions.

It is acknowledged that the CJEU has adopted the third and fourth stances in its rulings in cases 24/62 [195] and 162/97 [193] respectively. Over time, the proportion of recitals within directives has seen an upward trend. Kierkegaard [196] asserts that recitals enable member states to include normative elements representing a political compromise that could not be incorporated into the main text of the legislation, and that they enable the Commission to introduce normative elements without extending debates and disagreements. With regard to the nature of recitals in EU law, it is well known that the European Union is shaped by a purposive approach to drafting legal acts, which encompasses explicitly stated purposes as well as general principles mostly codified in policies and treaties. Recitals are integral for interpreting the intent behind laws, and though often regarded as lacking legal weight, they are frequently utilised by the courts for that purpose. While the CJEU has stated that recitals cannot override the actual normative provisions of a law, they still play a significant role in legal interpretation and can determine the validity of an act, thus emphasising their importance in EC law [120]. It follows that while recitals should not be disregarded, their exact impact remains ambiguous.

3.1.2 Legal text analysis and machine learning

A diverse array of NLP techniques have been applied effectively in the field of legal text analysis, showcasing the potent role of AI in this do-

main [197, 198]. The importance and usefulness of these techniques are well-documented in various comprehensive reviews [199, 200]. In practice, NLP frameworks are frequently used to facilitate machine learning tasks such as classifying judicial decisions [145, 201, 202, 203, 204]. Moreover, text mining has been instrumental for extracting structured data from legal documents [205].

In more recent developments, the application of deep learning (DL) models like neural networks has expanded in the domain of legal informatics, enhancing text classification, information extraction, and information retrieval [206]. Despite the sophistication of these newer models, it has been noted that traditional NLP bag-of-words methods such as TF-IDF and LDA often outperform more advanced context-aware techniques like BERT and Law2Vec in tasks like document similarity assessments [207]. According to this finding, traditional NLP methods are preferred for exploring co-occurrence networks within legislative texts.

The broader field of automatic legal text classification has seen various significant contributions. For instance, in [208], predictive models achieved considerable success in predicting US Supreme Court outcomes, with an accuracy of 69.7% for case outcomes and 70.9% for justice level vote outcomes over a sixty-year period. These models did not leverage network analytic features. Similarly, recent advances have demonstrated the efficacy of machine learning in predicting violations of the European Convention on Human Rights. Utilising vectorisation techniques with words (unigrams) or sequences of terms (bigrams or trigrams), they achieved an average accuracy of 0.75 and F-measures ranging from 0.62 to 0.85 [209].

Comparable achievements have been noted in semantic classification of norms in German laws, where machine learning approaches reached an accuracy of 0.83 [204]. These accuracy benchmarks provide a valuable comparative framework for evaluating the outcomes of classification experiments, even though they are applied in slightly different contexts.

3.1.3 Legal citation networks

Legal citation network analysis has evolved into a robust research field, characterised by diverse methodologies and addressing various specific challenges. The primary challenge involves identifying and extracting citations. Citations vary significantly in format, from comprehensive legislative citations to references to specific sub-articles [210, 211]. Techniques such as gazetteers, concept markers [212] and regular expressions [213, 211, 214] have been employed alongside more sophisticated methods like conditional random fields and bidirectional long short-term memory (BiLSTM) neural networks, which have shown comparable outcomes in identifying such citations [215, 210, 216]. A mixed named entity recognition approach that combines rule-based systems and supervised learning has also been applied to improve accuracy in citation detection and to correct errors related to typos and imprecise entities [217].

The second area of focus within this field is ranking legal documents. Network analyses often employ undirected edges between nodes, where edges are weighted by citation frequency, as seen in research on French legal codes [218]. The ranking of case law typically incorporates ways to measure the authoritativeness of, and implicit relationships between, judgments or parts thereof in wide-ranging judicial corpora, with some studies critiquing the use of in-degree centrality as an inadequate way to measure a case's importance [219]. Domain-specific methodologies may also consider variables like the court's authority, whether information was disseminated through official court websites, the presence of the judgment in jurisprudence publications, and its age [220, 221]. A similar network analysis approach has been taken with French legal codes [222].

Thirdly, the labelling of citations within these networks is addressed less frequently. Zhang et al. [187] assist legal practitioners by simplifying the citation gathering process through a recursive forward and backward chaining method; they identify the primary reasons behind each citation using a semantic-based approach. Meanwhile, a comprehensive

multi-relationship model of citation networks distinguishes between different types of references, such as amending or non-amending, and classifies sub-networks by legal basis, instruments cited, and impacts on legislation, among other criteria [190]. Sadeghian’s [189] approach focuses on extracting predicates around citations to classify relationships within legislation without referring directly to the content, using techniques like conditional random fields and k-means classification with word embeddings for edge classification.

Previous studies have also focused on pattern-matching to identify semantic relationships between legal concepts, constructing a typology of relations that has been foundational for further analysis [223]. While Maxwell et al. [224] explored explicit cross-references to external legislative texts, the research described in this chapter examines internal cross-references. Unlike most studies that prioritise explicit citations, Panagis et al. [225] delved into implicit citations by employing text similarity techniques to reveal deeper connections within texts, enriching our understanding of citation dynamics.

In line with the following case studies, judicial citation networks form a crucial subset of legal citation analysis and is primarily concerned with intertextual references in cases, rulings, or judicial decisions. The seminal study by Fowler et al. [226], which examined the implications of case law citations within the United States Supreme Court, significantly highlighted the utility of citation network analysis for understanding the doctrine of *stare decisis* and the broader application of legal precedents. This approach has been mirrored in continental legal systems to examine the judgments of high courts such as the Dutch Supreme Court [227], the Court of Justice of the European Union (CJEU) [228], and the European Court of Human Rights [229].

Langone’s [230] research demonstrates that conditional random fields and neural network models perform comparably in citation extraction, both identifying a vast array of references within legislative texts effectively. While methods like regular expressions are somewhat effective,

they often fail to capture every citation, and in particular, they often fail to distinguish between explicit and implicit references. In contrast, detailed, annotated legal judgments from Curia, the CJEU’s official website, ensures a more exhaustive identification of judicial citations.

The analysis of how legal sources are ranked based on citation data is also an integral part of this research area, often focusing on the interconnectivity of legal codes [218]. As mentioned previously, Derlén and Lindholm [219] have critiqued the reliability of in-degree centrality as a metric for evaluating judicial decisions, noting that a decision’s importance may increase if it is cited in subsequent pivotal cases. Moreover, domain-specific approaches consider various attributes, including the hierarchy of judicial authority [220].

The task of identifying and categorising citations also encompasses examining the text surrounding citations and the use of specific linguistic signals and patterns [231, 232]. Rather than solely ranking judgments, Sadl et al. [233] advocate ranking individual cited paragraphs to circumvent potential inaccuracies in the classification results, specifically within the complex citation system of the CJEU. This refined approach allows for the creation of a precise and detailed dataset. Besides this, network analysis has been effective in providing empirical evidence of traditional legal doctrines [234] including, for example, the role of *effet utile* in preserving the continuity and authority of European Union law [235], selecting the cases that mostly shaped the EU [236], and more general considerations on the nature of EU case law [237].

Recent network analysis tools have been used to enhance support for legal practitioners in their legal research. One such example is Justeus, a new legal information retrieval system for EU and Hungarian legal resources developed by Görög and Weisz [238]. Similar initiatives by Kuppevelt and van Dijck [239] have led to prototyping similar tools in the Netherlands, while other researchers have introduced citation recommendation tools to streamline the drafting process [240].

The uniform XML (eXtensible Markup Language) schema used by the European Union for case law documentation facilitates extending these methodologies across various European legal domains. One of the approaches outlined in this chapter incorporates recursive regression techniques to identify groundbreaking interpretative statements, thereby enabling a comprehensive historical examination of case law.

3.2 Network analysis in EU health law

The rapid expansion of computer technologies and systems is evident across various domains, including the legal sector [198]. Recently, there has been significant interest in automated extraction of information from legal databases and texts. This challenging task has led to the development of several approaches and tools designed to address the complexities involved [205].

The research described in this chapter concerns the nuanced topic of legal interpretation, which is the process of ascribing meaning to legal documents. The primary objective of this interpretive process is to discern the normative messages inherent in a specific piece of legislation by examining the relationships between different parts of that legislation. Identifying these interconnections is central to the hermeneutic process, particularly for the contextual interpretative method. The automated extraction of such information is a significant challenge, particularly analysing semantic interrelationships between parts of legislative texts. The next sections propose a methodological framework for developing an automated system capable of processing both explicit and implicit links within legislative texts. This necessitates thorough understanding of the meanings conveyed by the language used in these texts. The approach typically involves the use of annotated corpora and related analytical tools. Furthermore, the application of network analysis is explored by conceptualising interrelationships as a graph.

The following sections will detail the corpus-based and graph-based approaches used. Initial results will then be outlined before concluding with opportunities for future work.

3.2.1 Methodology

The framework employs a two-step approach. The first step is a corpus-based method to identify various types of interrelationships within a norm. This step includes establishing guidelines to facilitate human annotation. The outcome of this process is the identification of classes or labels, resulting in the creation of a gold standard corpus [241] that can be used with supervised learning algorithms in machine learning experiments [242]. The use of card sorting to identify patterns within the data will be considered. The classification effort may incorporate features derived from a traditional NLP pipeline.

The second step explores graph-based analysis to gain deeper understanding of the interrelationships within legislation and to enhance the feature set for classification tasks. Several metrics from network analysis are compared with the annotation output. The research question is whether graph metrics are useful for investigating interrelationships in norms within a general NLP framework for processing legislative texts.

This study builds on ongoing efforts to develop NLP resources and tools for legislative texts [185]. The case study concerns EU legislation, with the specific aim of mapping related recitals and articles.

A corpus-based approach

The process of annotating relations in legal text includes defining a schema and applying it to the legal document. The aim of the schema is to clearly define the kind of information that must be annotated. This phase includes an inventory of semantic categories to be used and specifying the granularity of the annotation. This step relies on the effort of experts in

the specific domain of the documents under scrutiny.

Annotation phase. At this stage, it is essential that each annotator works independently. This way, there will be evidence of any situations of disagreement, which will be resolved later. As a first step, the text of the corpus can be presented in a spreadsheet in order to facilitate the work of annotators. For instance, legal text can be split into different structural parts such as articles or paragraphs for comparison; and different Excel sheets can be used to separate different areas of comparison. It can be useful to ask annotators to keep track of their start and end time in order to calculate the average time required for the task at hand. Interannotation agreement can be measured using metrics such as Cohen's kappa [243]. Another annotator (adjudicator) can be involved later on to solve cases of disagreement.

Card sorting. Commonly used by information architects, card sorting is a popular user-centred method aimed at identifying patterns in data [244]. Participants, who are asked to work independently, group physical or digital cards, each displaying a piece of information, based on their own mental model of the information domain. More specifically, in closed card sorting, participants are provided with a set of initial groups; in open card sorting, they have no guidance, i.e., they define the groups they feel are most appropriate and then they provide a descriptive label for each group. Groupings produced by different participants can then be merged by means of a *card * group* matrix, where $cell_{i,y}$ shows the percentage of participants who assigned $card_i$ to $group_y$. In order to arrive at a shared classification, the general approach is to label each card with the group label that obtained the highest level of agreement among the participants. In the case of legal text, cards can be used to display single paragraphs. The output of this process is a series of norm groups, as in [179].

Graph-based NLP approach

The following paragraphs illustrate the various steps of the graph-based NLP pipeline adopted in the proposed methodology. The methodology begins with a traditional NLP pipeline followed by a network analysis approach.

Natural language processing. A traditional NLP pipeline is employed to preprocess legislative texts and obtain the stems of terms. Typical preprocessing steps include converting text to lower case and removing punctuation marks and stopwords.

Following the preprocessing, the terms are separated into tokens, which are then subjected to stemming. Stemming, in linguistics, refers to reducing words to their root forms. This enables the merging of singular and plural forms of the same term occurring in different sections of the legal text. Further analysis includes detecting parts of speech such as verbs, nouns, adverbs, and adjectives. This text processing phase can be implemented using established NLP libraries in common programming languages like Python or R.

Identifying features. The second step of the framework focuses on automated extraction of features for use in machine learning experiments. With a bag-of-words model, various features can represent the text. Beyond traditional n-gram models (sequences of n words), additional features include the frequency of different parts of speech, sentiment analysis of the words (e.g. categorising them as ‘positive’ or ‘negative’ according to dictionaries created for that purpose), word and character lengths, and term frequency (using TF-IDF).

Network analysis. Network analysis is conducted on two types of graph. The first graph connects norm types based on term co-occurrence. If the stem of a word is identified in different parts (e.g. recital 2 and Article 6) of a legislative text, an edge is created between these parts, with the edge weight representing the total number of co-occurrences. This graph illustrates interrelationships between norms from a linguistic perspective.

The second graph examines the role of terms within the document, linking different stems that co-occur within the same legislative text portion identified as norm type. The frequency of their co-occurrence determines the edge weight. For example, an edge between stem X and stem Y with a weight of 3 indicates three co-occurrences within the same norm type. Social network metrics from this graph, such as degree, betweenness centrality [245], and versatility [246], can be utilised as individual features for classification purposes. These metrics describe the role of each vertex in the graph in terms of its relationships with other vertices.

Weighted multi-layer networks. Further research on this topic can benefit from analysing multidimensional or multilayer networks, which have garnered increasing attention recently. Identifying interrelationships in legislative texts can enhance classification efforts, as different labels represent different relationship types, forming distinct network dimensions. Each label denotes a different layer, applicable to networks of stems and norm types respectively. Recent studies have introduced frameworks for investigating such complex networks with multidimensionality, proposing a set of basic concepts and analytical measures for these networks [247]. Specific metrics can be adopted such as the percentage of vertices or edges belonging solely to a specific dimension (dimension connectivity) [247].

Classification. This phase involves establishing an experimental setting for a classification task. The existence of a relation between two parts of the document is explored through a supervised machine learning experiment using a model trained with the annotation results. The framework encompasses several binary classification tasks where classes (labels) include: i) the existence of any relation, and ii) the existence of a specific relation type. Various classification algorithms, such as naïve Bayes, logistic regression, decision trees, and support vector machines, can be employed. For instance, a binary support vector machine can be trained with the labelled relationships from the gold standard corpus.

Evaluation. Classifier performance is evaluated by computing the well-established F-measure, which represents the ratio between precision and

recall. Cross-validation techniques, such as ten-fold cross-validation, are applied. In ten-fold cross-validation, the following steps are iterated ten times: i) the training data is divided into ten equal partitions; ii) the learning algorithm is applied to nine parts while testing on the remaining part. The final performance measure is the average of the ten parts. Once an estimation of model performance is obtained, the pipeline can be applied to new data.

3.2.2 Initial results

The following subsections outline the results of the graph-based NLP pipeline implemented with the methodology.

Annotation guidelines

Initially, two legal experts identified eight types of relationships between recitals and (sub-)articles based on their analysis of Directive 2004/23/EC in the Italian language. These relationships are as follows: Conceptually Similar (whether using the same or different wording), Constitutive (linking norms containing definitions of legal terms to norms containing those terms), Motivation (where one norm provides the principle or goal that motivates another norm), Impact (in terms of conflicting goals that may restrict one or both norms, or norms for enforcement or monitoring that impact the efficacy of classic deontic norms), Indirect Internal (norms A and C from the same legislation are linked indirectly where norm A cites another norm B from the same legislation, which is related in another way to norm C), Via Other Law (a norm related to another norm that cites another law and cannot be understood without reference to that law), Procedural (linking a norm describing a procedure by an EU institution to support the goal of another norm), and Contextual (linking deontic or other norms to norms that provide contextual information such as jurisdiction and entry into force), and Norm Group (where two or more norms are connected due to being part of the same general requirement). For

detailed descriptions and examples, please refer to [179]. Annotators were instructed to apply these labels to Regulation 141/2000.

A significant initial outcome of the corpus-based approach is the creation of a document that includes the annotation scheme provided by a domain expert. This document offers clear definitions for each type of relationship along with corresponding examples. These guidelines serve as standardised instructions for independent annotators¹.

Annotation agreement

Two annotators participated in this initial effort. Each annotator received a spreadsheet file with spaces at the top of each sheet to record the start and end times of the adjudication activity and the number of comparisons made. The annotators were asked to log the time spent during multiple work sessions.

To evaluate the annotation phase, the interannotator agreement was computed on the relationships that the legal experts had identified between three recitals selected from the legislation under scrutiny and the whole corpus of articles (or sub-articles where present) from the legislation. Among the pool of recitals, the first, the shortest, and the longest recitals were selected. Annotators had to determine the type of relationship (if any) between, for instance, the first recital (R1) and article 1, paragraph 1 (A1.1). If a relationship existed, a value of 1 was assigned; otherwise, a value of 0 was assigned. The results were expressed as a sequence of comma-separated values, listing triplets for each type: {R1,A1.1,0; R1,A1.2,1; R1,A1.3,0; etc.}.

In analysing the results, the interannotator agreement on the existence of any type of link between recitals and (sub-)articles (see Table 3.1) was examined first. The findings were as follows: 86 cases of agreement on whether a link exists or not, 54 cases of agreement on the presence of a

¹The document (in Italian) is accessible at the following link: www.di.unito.it/~sulis/NLPxLAW/LineeGuida_ClassificazioneTipiLegami.pdf

relationship, 32 cases of agreement on the absence of a relationship, and 28 cases of disagreement — 22 cases where only the first annotator identified a relationship, and 6 cases where only the second annotator identified a relationship. This resulted in an agreement percentage of 75.4% and a Cohen’s kappa agreement of 0.5. Despite this encouraging general agreement, significant disagreements were observed on the labels assigned to these relationships, with varied distributions among the classes (see Table 3.2).

While certain classes were identified with a relatively similar frequency, others show considerable discrepancies in their allocation. Qualitative analysis of the results suggests that some issues arise from differences between the directive on which the classification scheme was based and the regulation that was annotated. For instance, significant disagreement was found in the allocation of the Impact label, as indicated in Table 3.2. One annotator extended the Impact category to include the potential effects of planned future guidelines (based on consultations with member states and other parties) on the interpretation and effectiveness of the regulation’s stated goals. The annotation of further legislative sources will enhance understanding of such issues and allow the classification scheme to be fine-tuned. A more anticipated source of difficulty was the Via Other Law link, which involves determining whether it is necessary to consult the cited legislation to fully comprehend the article or recital in question.

The data reveals that the most common type of link identified was Motivation, which aligns with the recognised purpose of recitals — providing explanatory notes to facilitate understanding of the substantive provisions.

The results of this phase clearly highlight the complexity of the task and the need to refine the labels and revise the annotation guidelines. It is also evident that merely providing annotators with annotation guidelines is insufficient to adequately train them in a novel classification scheme. For future annotation exercises, seminars will be conducted covering both typical and challenging cases to ensure thorough understanding of the annotation task.

The annotation scheme involved documenting the start and end times of each annotation session. The average duration for each annotation session was approximately 62 minutes. With respect to the above-mentioned three recitals, one annotator completed the task more quickly than the other but divided the work into multiple shorter sessions. Specifically, this annotator conducted 6 sessions with an average duration of 24 minutes each compared to the other annotator who completed the work in 3 sessions with an average duration of 99 minutes each.

Grouping norms with card sorting

The card sorting exercise was conducted with the same individuals who performed the annotations. Although this task is still ongoing, preliminary insights from the material preparation phase suggest a need for methodological adjustments compared to traditional card sorting. First, the legal experts consulted for this activity noted that individual parts of legal text could belong to multiple norm groups. Consequently, multiple cards for each recital were prepared. Additionally, the number of cards generated exceeded the typical quantity used in card sorting, and the content on each card was more extensive. To address this, the annotators' cognitive load will be evaluated using both quantitative measures such as NASA-TLX (National Aeronautics and Space Administration - Task Load Index) [248] and in-depth qualitative interviews, in order to adapt the card sorting methodology to new application areas beyond information architecture.

Graph analysis

To process the legal text, standard techniques were applied to extract relevant information. For instance, performing POS tagging on the EU regulation yielded 1,636 unique words, with the most common type being nouns (1,239 occurrences), adjectives (171), and adverbs (58). Beyond typical NLP analysis, the document —recitals and (sub-)articles— was reduced by converting the words into their stems. The Python NLTK

(Natural Language Toolkit) library [249] was used for NLP processing, and the results were stored in a MySQL database to ultimately generate the edge list. The Gephi [250] open-source software was used to explore and manipulate networks. Two types of graphs were created to investigate interrelationships between norm types, as shown in Figure 4.9. To enhance the visualisation, vertices with low degrees were pruned in both graphs. The vertex size corresponds to degree, while the label size indicates betweenness centrality.

Graph of recitals and articles. The first graph examines the relationships between recitals and articles, with each relation represented as a vertex. This graph is depicted on the left in Figure 3.1 (a), including only vertices with a degree higher than 10 to improve network readability. Edges are weighted by the number of co-occurring terms in the two corresponding vertices, representing the strength of their relationship. Graph metrics indicate the strength of these relationships. Once the annotation of implicit links is complete, it would be interesting to correlate the interrelationships within a legal text with corresponding graph metrics, such as degree or centrality measures. It is hypothesised that there is a similarity between lexical graph-based measures and certain types of implicit relationships (e.g. Conceptually Similar, Motivation) within the text.

Graph of terms. The second graph analysis explores the relationships among terms after stemming. In this graph, stems are represented as vertices, and edges denote term co-occurrences within the legal text. Figure 3.1 (b) illustrates this type of graph, showing only vertices with a degree higher than 100. The edge weight represents the number of co-occurrences of stems in the document. Metrics for each vertex, such as degree and centrality, are calculated to measure relevance within the graph. Degree values range from 1 to 68, with an average of 52. Stems with higher betweenness centrality are semantically coherent with both the specific topic of the legislative act (e.g. “Medicinal” at 21.8; “Pazient” at 16.9) and to broader, general language (e.g., “Tal” at 21.9; “Scop” at 17.1).

Feature set

Each part of the legal text can be represented using a bag-of-words approach, including n-grams to capture more context around each term, as well as POS-based features (e.g. the number of nouns, adjectives, etc. in each part). The output of this step consists of a set of features collected as a numeric feature vector, aiming to create a feature model for binary classification. Information derived from graph analysis supplements the traditional feature model. Considering the graph with stems as vertices, each vertex metric can be used to create a feature corresponding to parts of the legal text, i.e. (sub-)articles or recitals. For instance, the average value of vertex metrics (e.g. degree, betweenness centrality, closeness centrality) can be utilised in the classification step. Similar considerations apply to the graph of relationships between stems co-occurring in different parts of the legal text. The final set of features may or may not include network metrics, depending on the network topology.

3.3 Co-occurrence networks of recitals and (sub-)articles in EU legislation

Recent years has seen a surge in legal informatics research, including in the domain of legal relationships extraction [251] and legal citation networks. Computational approaches in this area have advanced through the use of NLP and ML, along with an increased focus on graph analysis. Legal norms generally contain two kinds of relationships with other legal norms: explicit or implicit. Most research has focused on explicit citations to legislation or case law, whether expressed fully or partially.

This chapter focuses on the identification of implicit interrelationships within legislative acts, which are semantic relationships between parts of the text that do not explicitly refer to one other. Identifying these relationships is a significant challenge for both humans and machines, but it

is crucial for legal interpretation.

Automated systems are increasingly useful in supporting legal practitioners with the interpretation process [252, 253]. Various tools equipped with artificial intelligence (AI) systems can extract information from legal norms [198, 254]. This thesis proposes a general framework for the automated identification and classification of implicit interrelationships between parts of a legal text. The methodology combines a pipeline that includes NLP, ML, and a graph representation of the legal text.

A significant contribution of the approach is the adoption of *co-occurrence network (CN)* analysis. A textual CN is a network composed of items (unigrams, or words) represented as vertices whose importance or function can be computed through specific graph measures. This study demonstrates how graph metrics can improve the accuracy of the results.

In summary, this study addresses the following research questions:

- RQ1: How effective are interrelationship types in identifying implicit links between the recitals and (sub-)articles of legislative acts?
- RQ2: Can classification algorithms detect such links?
- RQ3: Can co-occurrence network analysis metrics improve the performance of classification algorithms?

Previous research has addressed the detection, resolution, and labelling of citations in the legal domain. However, there has been no systematic approach that exploits CNs to enhance legal classification efforts.

The task of identifying types of interrelationships between parts of legislation is non-trivial, and no annotated datasets currently exist. Therefore, it is necessary to create a manually annotated training dataset. This framework is aimed at the development of a system for identifying *implicitly* related norms, initially by human annotators and subsequently through automated means. The methodological steps outlined in the next section were applied in a practical use case involving the EU Regulation

on Medicinal Products [255]. The regulation is widely referenced in EU legislation and addresses healthcare, though it also concerns significant economic issues. Eleven years after its adoption, the regulation has been subjected to only minor amendments and remains largely in force. It was also chosen for its comprehensibility and length, being sufficiently long to provide a variety of implicit links but short enough to be manually annotated.

The legal text includes 11 recitals and 11 articles, and most of the articles are further divided into paragraphs, resulting in 38 (sub-)articles. Potential links between all recitals and all (sub-)articles are considered. Potential links between recitals and other recitals and between (sub-)articles and other (sub-)articles are not considered. Accordingly, the full combination of interrelationships between recitals and (sub-)articles explored in this work amounts to 418. Each combination can be annotated with one of eight labels, according to the typology described earlier.

3.3.1 Methodological framework

The analytical approach is structured into three steps. The initial phase concentrates on classifying types of norms links in a legal document. This challenging endeavour requires expertise from the legal domain. The approach involves a card sorting method and manual annotation of the corpus to develop a training dataset. The level of consistency among annotators will be used to evaluate the proposed link types between norms, responding to RQ1 concerning the efficacy of these link types.

The second step involves employing a graph-based methodology for feature extraction to augment the conventional NLP representation of legislative texts. The third and last phase investigates the impact of lexical features and network analysis metrics through several binary classification tasks, aiming to detect interrelationships within legislative texts. The outcomes of this automated process will address RQ2, which examines the classification capability of these relationship types, and RQ3, which




Figure 3.2: A comprehensive methodological framework for investigating interrelationships among norms.

assesses the value of network co-occurrence metrics.

Figure 3.2 illustrates the primary stages of the methodology, which is elaborated further below.

Identification of norms and the annotation process

Developing an annotated corpus of relationships within legislative texts is essential for conducting the classification analyses. This effort hinges on the expertise of domain specialists and encompasses several phases: delineating class categories or tags, crafting comprehensive guidelines, performing manual annotations on the corpus, and calculating interannotator agreement. A card sorting method 3.2.1 was implemented, primarily to assess the recognisability of different norm types and secondarily to test an alternative annotation approach from a user-centric perspective.

In the context of legislative texts, each card represents a paragraph, with the outcome being categorised sets of norms. Annotators' perceptions of card sorting's utility and usability are evaluated through a survey incorporating open-ended and 5-point Likert scale queries.

The annotation framework is meticulously defined to specify exactly what information needs marking, detailing the labels for use and the level of detail required for annotations. It is crucial that annotators work independently to capture any discrepancies, which are later addressed. Annotations are recorded in a spreadsheet to streamline the process. Interannotator agreement is quantified using Cohen's kappa 243. To resolve any

disagreements, two additional annotators with expertise in legal informatics are involved, with their consensus establishing a ‘gold standard’ corpus resulting from collaborative annotation efforts.

3.3.2 Results

The analysis of interrelationships within portions of legislative texts is enhanced through the use of two types of graph representations. Connections between sections sharing a common stem — the basic form of a word — are modelled, and the co-occurrence of stems within the same section is analysed. By focusing on stems rather than complete words, this approach finds a higher frequency of co-occurrences, offering more semantics into the related parts of legislation.

Network metrics were evaluated to elucidate the significance of each node’s connections to other nodes in the graph. The study specifically focuses on two distinct types of undirected graphs (G1 and G2).

A graph with text segments as vertices . In the first graph model (G1), each segment of the legislation, including recitals and (sub-)articles, is treated as a vertex. These vertices are connected by an edge if at least one stem appears in both segments. The weight of an edge represents the number of stems that co-occur across these segments. For instance, if a stem appears in both recital 2 and Article 6, these segments are linked by an edge. The thickness of the arcs depends on the number of co-occurrences.

A graph with stems as vertices .

The second graph (G2) focuses on the relationships between stems, where each stem constitutes a node. Edges in this graph are weighted by the count of document parts where the connected stems appear in both parts. For example, an edge connecting stem X to stem Y with a weight of




Figure 3.3: Illustrative sub-graphs of the stem co-occurrence networks, with G1 on the left designated as “a” and G2 on the right designated as “b”.

3 means that these stems are found together in three different parts of the document. To better illustrate the networks used in the study, descriptions of two sample sub-networks are shown below.

Illustrative examples of the networks

Figure 3.3 presents two “toy examples” of smaller subgraphs for each of the network models considered in this study, G1 and G2.

The G1 subgraph (Figure 3.3, [a]) shows four well-connected vertices: recitals *R5*, *R10*, *R11*, and Article *A8(2)*. The edges’ weights correspond to the number of co-occurring stems. The edge from *Article 8(2)* to *recital 10* is highlighted to demonstrate the graph construction process. The subsequent sections will detail the co-occurring stems identified, using *recital 10* as an example, with specific stems listed in bold.

The stems of recital 10 are:

programm specif biomed quart programm quadr ricerc svilupp tecnolog 1994-1998 sovven-
zion ricerc terap malatt rar metodolog istitu programm celer svilupp medicinal orfan
inventar medicinal orfan **dispon** europ tal fond intes promuov collabor internazional
mater ricerc bas ricerc clinic malatt rar comun continu attribui ricerc malatt rar import
prioritar previst **quint** programm quadr ricerc svilupp tecnolog 1998-2002 present regol
defin quadr giurid consent tempest effett applic **result** tal ricerc

The stems of Article 8(2) are:

tal period può tuttav esser ridott anni scadenz **quint** anno **risult** medicinal question
conform criter articol **risult** fra altro bas dat **dispon** rend tal giustific manten esclus
merc tal fin stat membr inform agenz criter bas concess esclus merc potrebb esser rispett
segu ciò agenz avvi procedur defin articol sponsor forn agenz inform necessar riguard

It is evident that between the two sections of the legislative act under examination, there are three common stems: *risult*, *quint*, and *dispon*. Notably, *risult* is mentioned twice in A8.2. The edge weight is computed by summing these occurrences, resulting in a total edge weight of 4 between A8.2 and R10.

In the G2 subgraph (Figure 3.3, [b]), attention shifts to the relationship between “medicinal” and “sponsor”. The process begins by converting the original Italian legal text into its stem components. For better clarity, the stems from Article 5(2) below are presented, where the co-occurring stems “medicinal” and “sponsor” in bold require special attention:

sponsor domand dev esser corred inform document sequent nom ragion social indirizz
permanent **sponsor** princip attiv **medicinal** indic terapeut propost giustific relat osserv
criter articol paragraf nonc descrizione stat sviluppi compres indic previst

This same procedure is applied across all sections of the legal text (i.e. all recitals and articles), totalling the instances of co-occurrence for the stems *medicinal* and *sponsor*. It is observed that these two stems appear together four times in specific parts of the legislation: R9, A5.2, A5.12, A6.1. Consequently, the edge connecting the vertices *medicinal* and *sponsor* in the graph is assigned a weight of 4.

3.3.3 Classification

The classification task primarily entails statistical text categorisation to derive automated rules from manually annotated training documents. Distributed word representations [256] is implemented, which is a conventional method for encoding each segment of legal text as a vector representation.

Specifically, each segment, i.e. recital or (sub-)article, in a legislative act d is represented as $x_{(i)}$, capturing the presence of sequences of words (bigrams and trigrams). The entire document is then depicted as a feature vector $x = (x_{(1)}, \dots, x_{(p)})$ of length p . The bag-of-ngrams (BoN) approach, often paired with the term frequency-inverse document frequency (TF-IDF) weighting scheme [257], computes $\text{TF}(i, d)$ as the frequency of a term i in document d . Dimensionality reduction is performed by selecting only the features with the highest document frequency.

The classification task examines both a binary scenario — determining the presence or absence of a relationship between two document parts — and the assignment of a specific class from eight classes. Given the complexity of multi-class problems for machine learning, binary classifications were carried out for each category.

In these binary classifications, instances where a relationship exists were treated as positive examples. A training dataset of manually classified pairs z was represented as $S = \{(z_1, y_1), \dots, (z_n, y_n)\}$ where n was the number of training pairs and $y_i \in \{1, 0\}$ denoted the class label. The binary classification leveraged a supervised learning framework aimed at deriving a classification rule — a function mapping from the p -dimensional feature space to a one-dimensional class label.

Representation of legal text

This stage involves the vectorisation of each section of the legal text by automatically extracting relevant features for machine learning analysis. Before vectorisation, the preprocessing of documents involves converting text to lower case, removing punctuation and stopwords, and tokenising to separate terms.

A bag-of-ngrams (BoN) model was chosen, focusing on stems rather than full words to enhance semantic understanding. For example, it allows matching of different forms of a term (singular vs. plural) appearing in various parts of the text. To emphasise less frequent but significant stems,

TF-IDF was applied, scaling the count of stem n-grams in a document by the inverse frequency of documents containing those n-grams.

The n-gram types selected were bigrams and trigrams; feature sets were created from pairs and triplets of sequentially appearing stems. Ultimately, the top 200 significant stems for both bigrams and trigrams (referred to as Bigr200 and Trig200 feature sets) were selected. Initial tests with a broader set of stems (e.g. the top 1,000 bigrams) did not yield better outcomes, leading us to concentrate on Bigr200 and Trig200 for the experiments.

Network analysis metrics

A broad range of network metrics are outlined in Table 3.3, with concise definitions and corresponding acronyms, categorised into three distinct groups (NM1, NM2, NM3). The first group (NM1) encompasses vertex-level metrics in graph G1, including Degr, DegC, BetC, CloC, EigC, LoaC, ClCo, Cons. The second group (NM2) pertains to edge-specific metrics in the same graph, comprising a boolean attribute *isLink* (1 if a link exists, otherwise 0), *Weig* (the weight of the edge), along with EdBC, CFlo, Effi, Conn, Cliq, kCli, and Jacc. The third group (NM3) relates to G2, where the average values for stem-associated graph-based metrics for each section of the legal text are computed. These include average centrality metrics like degree (avDegC), betweenness (avBetC), closeness (avCloC), eigenvector (avEigC), and load centrality (avLoaC), plus average values for two critical vertex properties: clustering coefficient (avClCo), and constraint (avCons). Table 3.4 lists these metrics across the two graph types.

Experimental setup

The methodology incorporates a series of supervised binary classification experiments. Initially, the occurrence of at least one relationship between two sections of the legal text is ascertained. The aim is to predict the presence of specific labels. Each textual fragment is represented by a com-

posite of features derived from both a bag-of-ngrams (BoN) model and graph-based metrics. A broad range of classifiers were evaluated, including logistic regression (LR), decision tree (DT), support vector machines (SVM), and k-nearest neighbours (kNN) [265], along with naive Bayes (NB) and a dummy rule-based (Du) classifier.

Validation and evaluation. Standard k-fold cross-validation for model assessment are employed. This well-regarded method for estimating a classifier's error rate reserves a data segment exclusively for testing. Specifically, the ten-fold cross-validation procedure is utilised (see [3.2.1]).

This analysis was executed using Python and its associated libraries, including NLTK [249], scikit-learn [266], and NetworkX [267], with additional graph analysis conducted in Gephi [250].

3.3.4 The annotation process

This section outlines the methodology used to annotate connections between individual norms and groups of norms, and provides an analysis of the results from these annotations along with feedback obtained after discussions.

Annotation results

Table [3.5] provides the principal findings from the annotation phase. Examination of raw figures suggests a relatively low concordance, attributed largely to the differing labelling tendencies of the annotators: Annotator 2 marked 617 relationships, significantly more than the 260 identified by Annotator 1, leading to notable variances in some classes. Nevertheless, the proportionality within individual classes remained roughly consistent, particularly in frequently cited classes such as Motivation, Via Other Law, and Procedural links.

The initial agreements are somewhat promising: the annotators concurred on the existence of 141 relationships and the non-existence of 2,608

relationships. Cases of disagreement, totalling 595, were adjudicated by additional experts. The resultant corpus encompasses 464 links spanning all eight categories.

Specifically, there were 139 instances where both annotators acknowledged the presence of a relationship and 99 instances where both recognised the absence of a relationship. Among the 180 disputed instances, the first annotator exclusively recognised a relationship in 22 cases, while the second exclusively did so in only 6. This resulted in an overall agreement rate of 57% and a Cohen's kappa of 0.2 on the existence of relationships.

Concerning the identification of specific link types, agreement rates were generally higher, ranging from 46.6% to 99.0%, largely due to a strong consensus on norm pairs lacking a specific class of link. However, Cohen's kappa values remained moderate.

Analysis reveals the highest agreement in classes defined by clear, objective criteria and formulaic language, such as Contextual and Constitutional links. Conversely, the Via Other Law class, which involves identifying citations, showed lower agreement rates due to differing applications of the class by the annotators. One applied it only when necessary for understanding the norm, while the other used it systematically whenever a norm referenced another law. Results for the Conceptually Similar and Motivation classes were modest, reflecting their more subjective nature.

The less favourable results for the Impact and Procedural classes stemmed mainly from differences between the directive and regulation under review. Although both pieces of legislation concerned the health domain, their focus differed significantly. Additionally, norms involving EU institutions like the European Commission, Parliament, Council, and specialist committees were often categorised as Procedural rather than Deontic norms, reflecting their descriptive rather than prescriptive nature.

During a post-annotation discussion with the annotators, it emerged that both the European Agency for the Evaluation of Medicinal Products and the Committee for Orphan Medicinal Products were viewed as

entities under the EU rather than independent bodies. However, norms related to these entities could be categorised as Deontic, Procedural, or a combination of both, depending on the specific content of the norm rather than solely based on its intended recipient. This indicates that the application of the Procedural category to the regulation involved a higher level of subjectivity than initially anticipated.

The Impact class showed the poorest results, with no consensus on any pair. Differences in interpreting this class in the directive and regulation were influenced by their respective contexts and goals. While the directive involved balancing conflicting goals like transparency and confidentiality, this was less apparent in the regulation. The broader application of the Impact class by one annotator to include the effects of future guidelines also highlighted interpretation discrepancies.

These findings indicate areas where further refinement of classes is necessary, suggesting a need for further annotation across more diverse legislation and the involvement of additional annotators. Further experiments should be conducted in the future on the Impact class. It may need to be split or discarded altogether, considering its wide scope.

Discussion of annotation results

The post-discussion revealed confusion over the use of the “Procedural” label for both a norm type and a link type, prompting suggestions for renaming these classes. Despite this, the guidelines for defining link types were seen as highly coherent, scoring 4.5 out of 5 in a post-discussion survey. However, one annotator noted that while the definitions were comprehensive, they could lead to ambiguity when interpreting norms. The annotators differed on the utility of pre-annotation training: one rated it moderately necessary, while the other highly recommended it, particularly if it includes practical exercises to assist understanding of the definitions among annotators.

Manual classification of norm groups

A preliminary card sorting activity was carried out by the two legal annotators. They were assigned two tasks. First, they were asked to group norms into the following six categories:

- Objective (Ob): outlines the purpose behind the directive as a whole, or some parts of it.
- Constitutive (Cns): contains definitions of technical concepts.
- Deontic (De): specifying types of behaviour required or permitted. May be further classified as permission, obligation, prohibition etc.
- Scope (Sc): outlines the extent of the applicability or non-applicability of norms (or the entire legislation).
- Procedural (meta-norm) (Pr): refers to step-by-step processes for implementing the legislation e.g. get signatures, get agreement from the committee, get further signatures.
- Contextual (meta-norm) (Cnt): refers to time, space, addressee and the hierarchy of norms.

Secondly, annotators were asked to sort norms into as many groups as they preferred and to identify a suitable label for each of them, using as a guide the *norm group* concept described in [179]: “A link between norms that are connected due to being part of the same general requirement. The links between the norms may be conjunction, disjunction or sequence. Such norms may be paragraphs of the same article, or may occur in different provisions.” It is envisaged that the identification of such norm groups may not only be useful in itself but may also help to improve the identification and classification of links between norms.

The methodology employed for this exercise was card sorting (first task: closed; second task: open [244]). Each annotator was provided with a set

of physical cards, with each card displaying a single paragraph. Forty-nine different cards, corresponding to 11 recitals and 38 articles, were included in each set. Multiple copies of the same cards were available to allow the annotators to allocate these cards to multiple groups.

Note that while this activity is normally carried out with 15-20 participants [268], a preliminary card sorting activity was conducted with only two annotators in order to obtain some feedback on the difficulty of the task and the viability of card sorting as an alternative annotation method when labels are applied to sets of items.

Card sorting results

The following paragraphs summarize the results of both closed and open card sorting experiments, as well as a discussion of their findings.

Closed card sorting Table 3.6 details the outcomes of the closed card sorting exercise. A total of 115 card-to-category allocations were recorded, with 16 out of 50 cards (33%) placed in multiple categories by the annotators themselves. There was consensus on 20 of these matches, including 8 cards categorised uniformly by both annotators and 2 cards placed into the same two categories by both, resulting in complete concordance for 10 out of 50 cards (18%). Overall, a Cohen's kappa of 0.19 suggests minimal agreement.

There was notable variability in the distribution of cards across categories by the two annotators. The highest agreement was found within the Contextual (meta-norm) category, where both annotators categorised the norms identically (Cohen's kappa: 1), and the Objective category, where approximately half of the cards were similarly classified by both annotators (Cohen's kappa: 0.56). In contrast, there was no agreement with regard to the Scope and Procedural (meta-norm) categories, which reflected previous findings from the annotation of link pairs.

Open card sorting Table 3.7 presents the findings from the group open card sorting task. One annotator formed 6 groups while the other formed 16, with an average of 10 and 7 cards per group respectively. The structures and labels of these groups showcased the annotators’ different cognitive approaches; notably, agreement on labelling occurred in only two instances, specifically for the categories *Committee* (Cohen’s kappa: 0.40) and *Market Exclusivity/Monopoly* (Cohen’s kappa: 0.55), as detailed in Table 3.8. Overall, they agreed on the classification of only 7 card-norm groups.

Significant overlap was observed, with 41 out of 49 cards (84%) categorised into multiple groups by individual annotators, indicating a higher identification of norms in their corresponding norm groups as opposed to grouping based on norm types. One annotator explained that her classification into various groups and subgroups was intentional, reflecting her hierarchical perspective on norm groups.

Discussion on card sorting results In assessing the recognisability of norm types, the closed card sorting phase underscored that certain concepts, particularly *Scope* and *Procedural* (meta-norm), were interpreted differently by the annotators. These observations confirm previous findings on the complexity of classifying norms. Accordingly, it will be necessary in the future to have a reassessment of the norm classes and refinement of the annotation guidelines to enhance clarity and consistency.

Moreover, open card sorting appeared to be effective for tasks that anticipate a hierarchical arrangement of concepts and where labels cannot be predetermined. However, establishing a “gold standard” may be challenging as it requires managing disagreements not only within each norm group but also across different groups.

From a usability perspective, annotator feedback suggested that card sorting was perceived as less efficient and more cumbersome than spreadsheet-based methods (rated 3.5 vs. 4.5 out of 5). Despite some annotators finding the card sorting method somewhat disorganised, the physical han-

dling of cards was seen as beneficial. It was recommended that annotators should be provided with a printed overview of the legal text to better contextualise their card sorting tasks.

Given these insights, card sorting could offer multiple advantages if integrated into the annotation framework:

- It may serve as an initial step to gauge the consensus on labelling among annotators, prompting adjustments where necessary.
- When defining norm groups, it could act as an introductory training phase, enhancing annotators' comprehension of the legal text and their proficiency in identifying and categorising normative links.
- Assuming enhancements to the user experience, it could become a viable alternative to conventional annotation methods.

3.3.5 Graph analysis

Graph representations enhance the analysis of structural connections within the document manifested through lexical links. To characterise the two distinct graphs employed in this study, a range of different network metrics were applied. These include the number of vertices ($\#V$) and edges ($\#E$) as well as average degree (AvDegr), diameter (Dia), average path length (AvPaLe), density (Den), transitivity (Tra), and average clustering coefficient (AvClCo). Table 3.9 presents the key attributes of both Graph 1 (G1) and Graph 2 (G2).

The following sections detail the set of network features arising from the two graphs.

Graph of recitals and articles. To illustrate the structure of the graphs created, Graph 1 (G1) is presented in Fig. 3.4, which visualises the components of the text, recitals and (sub-)articles, as nodes. These nodes are interconnected wherever there is at least one common stem shared between them. The node sizes correlate with their degree, whereas the size




Figure 3.4: A representation of graph $G1$ concerning relations between recitals and (sub-)articles. The vertices are recitals (R) and (sub-)articles (A), where edges represent a link if two vertices have at least one common stem.

classification of each node is based on their betweenness centrality.

The thickness of the edges reflect their weight, which is determined by the number of stems shared between the paired nodes. Such weights gauge the robustness of the connections between different sections of the legislation.

To delve deeper into the network structure and uncover more tightly-knit communities within it, the Louvain modularity algorithm [269] was applied. This method distinguishes between four distinct communities within the graph, differentiated by colour, containing 18, 14, 12, and 5 nodes respectively². While comprehensive analysis of these groupings is beyond the scope of this thesis, it opens avenues for future exploration into the nature of the relationships depicted.

The network metrics provide some insights about the graph topologies.

²Access to the code and the colour-enhanced graph is available via a GitHub repository: <https://github.com/sulem76/LegNet>

G1 is a graph that is relatively compact, having a quite fair density (about 0.3). In fact, it is quite easy to reach most other vertices with a limited number of connections: the average shortest path length is 1.7, and the diameter of the whole network is 3. The average degree is significant: each part of the text is connected to 13 other vertices on average. Values of transitivity and clustering coefficient indicates an important connectivity, i.e. a certain local neighbourhood where a region is connected to its neighbours.

Graph of stems. G2 represents a relatively sparse graph, which aligns with expectations given the connections between stems. The graph exhibits low density, with the average degree showing that each of the 548 stems (represented as vertices) is connected to fewer than 5 other stems on average. Despite this, the graph does not spread extensively, as it has a diameter of 7 and an average path length just above 3. Unlike G1, which demonstrates the characteristics of a small-world network (evidenced by both a high clustering coefficient and a short average path length), the metrics of G2 reveal a limited number of connections among neighbouring stems.

In analysing G2, the metrics associated with each vertex (or stem) can serve as features related to specific segments of legislative texts like articles or recitals. Indeed, the average values of vertex metrics such as degree, betweenness centrality, and closeness centrality can be leveraged during the classification phase.

3.3.6 Classification outcome

The relationship between different segments of the legal document is investigated through a supervised machine learning experiment, using models trained on annotated data. The first group of experiments (S1) examines the presence of any type of interrelationship, regardless of its nature, between two parts of the legal text. In the second set (S2), the focus shifts to predicting specific individual labels. Finally, the third group of experi-

ments (S3) explores various configurations of the network-based features, with particular attention given to the role of the connections between different sections of the text.

Inferring interrelationships

The initial experiments demonstrate that the comprehensive set of features can predict the presence of a relationship between two segments of legal text with a notable degree of precision (an F-measure of approximately 0.84 using both logistic regression and support vector machines). The performance of two groups of features, Bigr200 and Trig200, are comparable (with trigrams performing slightly better than bigrams in four out of five classifiers).

Table 3.10 presents the performance of various classification algorithms, detailing both their F-measure and accuracy. Considering the complexity of the task and the baseline established by the dummy classifier, these results are quite good.

Inferring link norm categories

Predicting individual classes presents multiple problems. First, the types of norm links carry nuanced meanings that are not readily captured through lexical features or relationships. Additionally, some norm link types are infrequent, necessitating extensive annotation efforts to create a gold standard corpus that yields relevant results.

Table 3.11 provides binary classification outcomes for a selection of the most effective algorithms (logistic regression, decision trees, SVM) using bigrams (Bigr200), though results with trigrams are similar. The F-measure for the four most common norm link types in the annotated corpus are evaluated: MO, II, VOL, and PR. These results demonstrate the classifiers' ability to address the complex issue of identifying specific types of relationships.

The impact of network features

This section examines the effectiveness of using network metrics independently, without combining them with a bag-of-ngrams (BoN) representation. As shown in Table 3.12, the results indicate that network metrics alone yield an F-measure of approximately 0.81 with two classifiers. Although these values are slightly lower compared to those achieved with ngram-based methods alone (around 0.83), they highlight the benefit of integrating network metrics with a BoN model because it results in a higher F-measure of roughly 0.84. Despite the modest improvement in performance, the consistency achieved across the different classifiers employed in this study suggests that this approach is worth further investigation.

To clarify the influence of network features on classification, an analysis was conducted focusing on various subsets of network metrics. Specifically, the inclusion or exclusion of specific network metrics was adjusted, as outlined in Table 3.4, to assess their individual effects on classification tasks. This feature removal approach involved different configurations of the sub-groups NM1, NM2, and NM3 (described in Table 3.4).

Each configuration is represented by a combination of T (true) and F (false) values indicating whether a sub-group is included or excluded. For instance, TTT signifies that all three groups are included, TTF includes only NM1 and NM2, TFT includes NM1 and NM3, and so on.

Results are presented from using the Bigr200 feature set, and comparable results are obtained for Trig200. The outcomes of this experiment are illustrated in Figure 3.5. Two key patterns stand out. 1) The best F-measure performances were achieved with the LR and SVM algorithms, while DT demonstrated superior accuracy. 2) The configurations that produced optimal results were FFT, TFT, and TFF, consistently highlighting the importance of edge-based metrics (NM2), which were included in all the top-performing configurations. This suggests that further investigation into NM2's role in enhancing classification performance is warranted.




Figure 3.5: F-measure (radar plot on the left) and accuracy (right) for logistic regression, decision tree and support vector machines classification of network metrics configurations

3.4 Judicial citation networks: a CJEU case-study

According to prevailing legal doctrines, judgments in common law systems are typically binding on subsequent decisions, while in civil law systems, they serve primarily as persuasive references. This stark distinction, however, has been increasingly questioned. Some scholars contend that the actual difference in the application of precedents between civil and common law systems is minimal, suggesting that courts in both systems adhere to precedents to a comparable degree [270]. Additionally, the citation style of the Court of Justice of the European Union (CJEU) is notably distinct, blending civil law traditions with a common law approach to referencing its own precedents [271].

The importance of citing precedents within CJEU case law has grown significantly, as evidenced by an increasing trend in the number of citations per judgment over time. This shows the importance of systematically analysing these references to fully grasp the evolution of the CJEU’s judicial framework.

Traditional legal methodologies offer ways to analyse judicial decision-making through citation networks, but these manual approaches are often labour-intensive and susceptible to biases, particularly those stemming from non-empirical legal doctrines such as *opinio iuris*. A notable study by Derlen and Lindholm [219] investigated the citation network within CJEU case law and identified *Dassonville* (1974) and *Bosman* (1995) as key landmark judgments based on their citation frequency in subsequent CJEU decisions. This finding contrasted with that of legal textbooks which often highlight *Cassis de Dijon* (1979) and *Van Gend en Loos* (1963) as more foundational. This discrepancy underscores the potential of computational legal analysis to augment traditional legal studies by offering novel insights into the patterns of judicial decision-making.

This third part of the chapter investigates the citations of prior rulings by the Court of Justice of the European Union (CJEU) in appellate decisions concerning fiscal state aid, a domain notably shaped by CJEU judgments due to limited foundational legislative texts [272]. More precisely, appeals against decisions from the General Court (GC) were examined, focusing on this subset due to the relatively recent inception of these judgments (beginning in 1997), which ensures a more uniform citation style [233].

Additionally, the Publications Office of the European Union has implemented an XML schema across CJEU judgments, facilitating the automated extraction of citations using robust and reliable data, thereby enhancing the accuracy and reliability of the analysis.

The primary aim of this research is to delineate and visualise citations to precedents within CJEU case law on fiscal state aid. The data extracted provides a foundation for graphical representations to support legal scholarship that might enable advanced analytics through natural language processing (NLP) and artificial intelligence (AI) techniques [273]. The methodology leverages a corpus of CJEU cases, preferably in XML format. Where XML data is unavailable, regular expressions based on heuristic rules are used to meticulously extract references to other CJEU decisions.

The resulting data, stored in JavaScript Object Notation (JSON) and graph formats, is subsequently subjected to network analysis, potentially using tools such as Neo4j.

This initial phase sets the groundwork for more extensive investigations that bridge linguistics, legal theory, and computer science.

3.4.1 Methodology

This research specifically targeted a collection of judicial decisions from the fiscal state aid sector, accessible via the EUR-Lex website. Although the focus was narrow, the methodologies developed in this study are adaptable to a broad range of legal fields, contingent only on the citation styles utilised in the judgments.

The extraction process commenced with the retrieval of XML data from the EUR-Lex platform [274], which offers a systematic way to handle information for research purposes and thus facilitate preliminary analysis. The XML format, comprehensively described in [275], employs various tags to denote citations. For this research, the primary focus was on the *REF.DOC.ECR* tag, which provides exhaustive references to prior cases. Initially, attention was limited to specific paragraph references, which were considered to be the most pertinent legally, though this parameter can possibly be modified in future research expansions.

It should be noted that not all decisions are available in the XML format. Where XML files were unavailable, regular expressions were applied to capture relevant data in order to work with both structured and unstructured text formats. When parsing XML, the marked and verified data already present in the file proved invaluable.

In typical instances, references within CJEU cases appear as follows:

[...] see, inter alia, judgments of 15 November 2011, Commission and Spain v Government of Gibraltar and United Kingdom, C-106/09 P and C-107/09 P, EU:C:2011:732, paragraph

From this, a list of tuples was compiled that included both the case identifier, in adherence with ECLI (European Case Law Identifier), and the paragraph number.

Upon identifying a cited paragraph, the corresponding case document was downloaded and the pertinent text was extracted. This reapplication of the reference extraction method to individual paragraphs enabled us to selectively aggregate citations of relevance to the legal analysis. The results of this refined collection process were particularly evident in the differentiated clusters of legal precedents versus procedural rules in the citation network.

With recursive application of this method, a layered citation structure was created with each layer representing an element named *references* in the data structure.

Enhancing the parsing of the XML documents with regular expressions proved crucial for accurately determining paragraph numbers and content. This approach aligns with established practice in natural language processing for extracting structured information from plain text legal documents [214, 213]. The structured data was then converted into the JSON format to optimise their use by automated systems.

The code provided below illustrates how the JSON structure formats and stores information about each reference including various identifiers such as ECLI (European Case Law Identifier), CELEX (Communitatis Europae Lex), and the case number, as well as the judgment name and the text of the cited paragraph:

```
“ ‘ecli’: “ ‘ECLI:EU:C:2011:732” ,
“ ‘text”: “ ‘C-106/09 P and C-107/09 P Commission and Spain v Go
“ ‘par_num”: “ ‘NP0073” ,
“ ‘celex”: “ ‘62009CJ0106” ,
“ ‘case_no”: “ ‘C-106/09 P” ,
“ ‘xml_url”: <url to the xml representation> ,
```

```
“references”: [...],  
“par_text”: “On the other hand, advantages resulting from a
```

With this data, it is feasible to develop visualisation tools, citation maps and databases to enable advanced analysis of the citation network.

Neo4j was employed for importing and visualising the data, utilising graphical or tabular displays to elucidate the relationships within the citation network. Initially, the system constructs a tree structure based on the “references” data entries. This structure facilitates deeper analytical examination by illustrating the interconnections between various graph nodes, which represent cases and cited paragraphs. In the visual representation, cases are depicted as yellow nodes and paragraphs as purple nodes. Two primary relational types are distinguished: REFERS_TO, which indicates a citation link between paragraphs, and BELONGS_TO, signifying the association of a paragraph with a specific case.

The subsequent section will explore one of the methodologies for querying this database, enhancing the performance of legal analysis on judicial precedents.


3.4.2 Results

Building upon the data structure outlined previously, the system constructs graphical depictions of citation relationships. These visualisations facilitate understanding of the chronological progression of citations and the development of legal doctrines.

Figure [3.6](#) illustrates a path from an initial case, shown in yellow, through subsequent paragraphs cited from prior rulings, depicted in purple. Each cited paragraph provides some context about its source and the judicial decision it accompanies. This visual tool is invaluable for analysing the trajectory of a legal principle from its inception in jurisprudence to its status as an established doctrine [\[276\]](#).

The utility of this visual representation is further enriched by its ability

Figure 3.6: Visualisation of a citation branch originating from a specific case



to highlight both the text of the citations and their interconnections, helping legal scholars to evaluate the profound influence of CJEU jurisprudence on the European legal framework.


It is important to note that Figure 3.6 represents only a singular strand of the broader citation network. The visualisation can be expanded to include various queries that extract specific information relevant to particular cases. This allows a detailed view of relationships among clusters of cases or focused examination of sections of particular interest.

With this dataset, future research efforts can use advanced algorithms to analyse citations in terms of textual similarities and assess the usefulness of network nodes.

Figure 3.7 presents a segment of the database, illustrating the links between various cases, specifically highlighting instances where multiple cases reference the same paragraph directly (excluding recursive citations). This visualisation facilitates the application of algorithms like PageRank [277] to automatically assess relationships and pinpoint the most legally significant citations.

Currently, the database predominantly covers state aid cases, but it is structured to seamlessly incorporate a broader spectrum of judgments.

Figure 3.7: Interactions between multiple cases within the database



The principal challenge in broadening the dataset lies in the variable quality of the input data. A preliminary review on the EUR-Lex platform indicates that not all cases are available in XML format, and those that are sometimes lack completeness. However, as mentioned earlier, missing data such as paragraph numbers and identifiers can often be retrieved from the text version using regular expressions.

Another complication arises when citations lack a unique identifier like an ECLI or CELEX code. Often, judgments can be identified based on the parties involved. But sometimes there are multiple other cases involving the same parties, e.g. the European Commission versus a certain member state. Here, heuristic methods could prove beneficial, for example, focusing on precedents to the case under review and taking into account the tendency of more recent judgments to have well-structured citations.

Additionally, the occasional unavailability of judgments in English poses a minor issue, though it is generally resolved through machine translation and revision by legal experts. However, using official, original multilingual versions of judgments is better than resorting to automated machine translation.

Preliminary tests on other legal domains suggest that judgments related to fiscal state aid consistently present well-documented precedents, both in XML and natural text formats. This may be due to the absence of other legal instruments such as directives or regulations, thus elevating the role of case law in defining the legal framework of fiscal state aid.

3.5 Conclusions and future work

The chapter illustrates three case-studies in which different network analysis approaches have been adopted. In the first case-study, the data-centric framework outlined is aimed at automatically identifying implicit internal links between norms within the same legislation. Specifically, a corpus of annotated links that map recitals to (sub-)articles was developed.

The methodology was initially applied to EU legislative acts. The initial analysis conducted by two annotators demonstrates the complexity of the task: while there is a satisfactory level of agreement in identifying interrelated parts, there is significant disagreement regarding the types of relationships. In addition to the corpus-based approach, a classification task was conducted using a set of features derived from traditional NLP methods. The integration of features from graph analysis was also explored to enhance the results. In this context, graph-based metrics related to vertices could serve as valuable additional features in a machine learning experiment.

For future work, these metrics will be employed to perform binary classification tasks. The classes will correspond to those used in the annotation phase, such as the presence or absence of a link and the type of relationship. Special emphasis will be placed on multidimensional network analysis [278], where different networks of interrelationships can be constructed based on the annotation labels to uncover new features of interest.

In the second case-study, a comprehensive NLP framework designed

to automatically classify implicit connections between different segment in a piece of legislation was outlined. The links corresponding to various relationship types were identified through a detailed annotation process, carried out by two legal annotators, with any disagreements resolved by two legal informatics researchers. Initially, the annotators identified and classified one-to-one links between different parts of an EU directive. These annotated results were subsequently reviewed to create a refined “gold standard” corpus, which served as the basis for developing an automatic identification and classification system.

To address RQ1, the “card sorting” method was incorporated into the workflow, allowing annotators to group norms based on type and semantic similarity. This process provided valuable insights into link pair annotation and facilitated the refinement of both norm and link type categories. As discussed, while the link classes used in the annotation guidelines were considered coherent and well-defined by the annotators, some differences in interpretation arose in practice. For example, challenges were noted in interpreting the Impact and Procedural link classes. Even in more straightforward categories like Conceptually Similar and Motivation, disagreements and uncertainties were observed due to the inherent subjectivity involved in applying these classifications to norms, which varied in content and could be interpreted in multiple ways. Following the creation of the annotated corpus, a machine learning experiment was conducted using a graph-based NLP approach. The classification results were promising, indicating that implicit links and their respective classes could be automatically identified using standard algorithms, thereby addressing RQ2.

In addition to employing a BoN model based on the stems, it was demonstrated that network metrics (CN metrics) could further enhance the identification of certain implicit links between norms. As for RQ3, a slight improvement in accuracy across was observed almost all classifiers when incorporating the full set of features. For example, the top-performing classifier (SVM) achieved an F-measure of 0.838 with all features, com-

pared to 0.832 using only BoN and 0.808 using only network features. Notably, edge-related features (NM2) showed potential for further exploration in future studies. Lastly, the classification results are generally in line with related studies, and the observed interannotator disagreements reflect the inherent complexity of the task. Nevertheless, in future work, the classification framework will be enhanced, which should lead to improvements in the performance of classification algorithms. While the subjective nature of even well-defined categories will likely continue to allow some level of disagreement, an automated system, even if not perfect, could significantly enhance the efficiency legal professionals' search for implicitly related norms.

Another interesting opportunity for future research would be to complement this work with qualitative analysis. While this chapter offers insights into the feasibility of identifying various types of links, it is equally important to determine which types of links are most valuable to legal practitioners in order to better focus future efforts. To gather this feedback, practising legal professionals should be involved, for instance through a seminar and a detailed survey. Based on the insights gathered, the link type model will be refined, addressing the challenges identified in this chapter and incorporating input from legal experts. It would also be beneficial to involve additional annotators and apply the annotation process to different types of laws and links between recitals and (sub-)articles. Another key focus of future work will be exploring how drafting rules, as outlined in documents like [279], can help creating norm types from scratch.

With regard to the CJEU case-study, the visualisations offered by the system serve as powerful aids for European case law analysis. The vertical display (Figure 3.6) traces citation chains from a specific case, offering insights into the decision-making context and its influence on subsequent legal outcomes. This format could also be invaluable for legal professionals, facilitating the exploration and interpretation of legal principles and their development over time.

This preliminary study sets the stage for extensive future research pos-

sibilities. The citation network could benefit from advancements in NLP and ML to enhance semantic analysis of text surrounding citations by enabling the extraction of meaningful data and examining their impact on judicial decisions. The network could also be developed to highlight the most frequently cited nodes (paragraphs or principles) and track how their citation frequency evolves. Utilising the complete network can help pinpoint key paragraphs and enable comparison of the text in citing and cited paragraphs. This comparison could be instrumental in identifying newly established legal principles.

From a legal perspective, the network's utility could be augmented by integrating additional data extracted from legal cases. Possibilities include correlating citations with case outcomes and extracting relevant keywords to classify cases and their citations by procedural type, specific issue, originating state, or legal impact. Moreover, the preliminary results could be refined using automated NLP tools to extract further information from cited paragraphs, eliminating the need to parse entire texts. This could help determine the extent of similarity across cited paragraphs from different judgments to ensure consistency in the meaning conveyed.

Another approach could involve associating citations with keywords extracted in all available languages, possibly organised using NLP tools and other automated methods. This would facilitate understanding of how keywords are handled across different European languages and help to develop a unified, multilingual ontology for each pertinent sector. Initially, this work leverages the XML format used by the CJEU itself to ensure that the findings are free from interpretative bias and that extraction is swift without requiring manual annotation. Consequently, a tool that enables anyone to access a decision's citation network by simply inputting its ECLI can be provided, making the output immediately useful even without the need for further NLP technology applications.

Table 3.1: Annotation results on the existence or non-existence of an interrelationship.

	Recital 1	Recital 4	Recital 8	All 3 recitals
Number of recital-article pairs	38	38	38	114
Annotator 1: number of links (of any type)	21	36	19	76
Annotator 2: number of links (of any type)	24	25	11	60
Agreement on the existence of a link	20	23	11	54
Agreed cases of a lack of link	13	0	19	32
Agreement on whether a link exists	33	23	30	86
Disagreement on whether a link exists	5	15	8	28
Link assigned by Annotator 1 only	1	13	8	22
Link assigned by Annotator 2 only	4	2	0	6

Table 3.2: Annotation results concerning different types of implicit links identified between three recitals and the whole set of normative provisions from Directive 2004/23/EC. CS = Conceptually Similar, CNS = Constitutive, MOT = Motivation, IMP = Impact, II = Indirect Internal, VOL = Via Other Law, PR = Procedural, CNT = Contextual (see Section 3.2.3 for full descriptions).

	CS	CNS	MOT	IMP	II	VOL	PR	CNT
Recital 1 links: 1st Annotator	0	3	16	1	4	7	6	1
Recital 1 links: 2nd Annotator	0	2	4	9	4	8	8	1
Agreement on Recital 1 links	0	1	2	0	2	6	5	1
Recital 4 links: 1st Annotator	1	3	30	0	5	8	5	1
Recital 4 links: 2nd Annotator	3	2	14	10	3	7	8	1
Agreement on Recital 4 links	1	2	9	0	3	5	2	1
Recital 8 links: 1st Annotator	0	3	16	0	5	19	5	1
Recital 8 links: 2nd Annotator	5	1	8	0	0	3	3	1
Agreement on Recital 8 links	0	1	8	0	0	3	2	1
Links for all 3 recitals: 1st Annotator	1	9	62	1	14	34	16	3
Links for all 3 recitals: 2nd Annotator	8	5	21	19	7	18	19	3
Agreement on links for 3 recitals	1	4	19	0	5	14	9	3

Table 3.3: Description of network metrics, with their respective acronyms in brackets

Name	Description
Degree	The number of edges that are incident to a vertex (Degr).
Clustering Coefficient	The fraction of possible triangles through that vertex (Clus) [258].
Centrality	Centrality metrics [245] address the position of the vertex in a graph by considering different perspectives, including some e.g. the neighbours of the vertex: degree centrality (DegC); the importance of its position in the graph: betweenness centrality (BetC); the distances from other vertices: closeness centrality (CloC); the influence of the vertex in the network: eigenvector centrality (EigC); the influence of a node in a network based on the influence of its neighbors, meaning that a node is considered important if it is connected to many other important nodes; current-flow betweenness centrality (CFlo) [259]. Edge betweenness centrality (EdBe) is also considered and, as its name suggests, focuses on edges.
Constraint	A measure of the extent to which a vertex is invested in those vertices that are themselves invested in the neighbours of the vertex (Cons) [260].
Clique	A subset of the vertices of an undirected graph such that each two distinct vertex is adjacent. Vertices and edges of the maximal clique (Cliq) are distinguished, as well as k-clique communities (kCli) [261].
Jaccard Coefficient	The Jaccard similarity index (Jacc) computed between all pairs of vertices [262].
Efficiency	The efficiency of a pair of vertices computed as the multiplicative inverse of the shortest path distance between the vertices (Effi) [263].
Connectivity	The local edge connectivity, which is the minimum number of edges that must be removed to disconnect them (Conn) [264].

Table 3.4: Network metrics sets (NM1, NM2, NM3) derived from G1 and G2 graphs

Name	Graph	Network Metrics
NM1	G1	Degr, DegC, BetC, CloC, EigC, LoaC, ClCo, Cons
NM2	G1	isLink, Weig, EdBC, CFlo, Effi, Conn, Cliq, kCli, Jacc
NM3	G2	avDegC, avBetC, avCloC, avEigC, avLoaC, avClCo, avCons

Table 3.5: Annotation results from Annotator 1 and Annotator 2, with cases of initial agreement and disagreement. The final annotated corpus after resolution of disagreements (by an adjudicator).

	CS	Co	Mo	Im	Pr	Cx	II	VOL	Total
Annotator 1	15	16	81	38	53	11	11	35	260
Annotator 2	12	23	254	6	122	13	56	131	617
%Annotator 1	5.8	6.2	31.2	14.6	20.4	4.2	4.2	13.5	100
%Annotator 2	1.9	3.7	41.2	1.0	19.8	2.1	9.1	21.2	100
Agreement YES	3	10	56	0	30	10	6	26	141
Agreement NO	394	389	139	374	273	404	357	278	2,608
Disagreement YES	21	19	223	44	115	4	55	114	595
YES - Annotator 1 only	12	6	25	38	23	1	5	9	119
YES - Annotator 2 only	9	13	198	6	92	3	50	105	476
Percentage of agreement	95.0	95.5	46.6	89.5	72.5	99.0	86.8	72.7	82.2
Cohen's kappa	0.2	0.49	0.06	n/a	0.20	0.82	0.14	0.20	0.24
Gold standard corpus	9	18	238	3	57	11	55	73	464

Table 3.6: Closed card sorting results from two annotators (Annotator 1 and Annotator 2), with cases of agreement and disagreement. Ob = Objective, Cns = Constitutive, De = Deontic, Sc = Scope, Pr = Procedural (meta-norm), Cnt = Contextual (meta-norm).

	Ob	Cns	De	Sc	Pr	Cnt	Total
Annotator 1	5	11	31	2	0	1	50
Annotator 2	11	3	16	0	34	1	65
Both	16	14	47	2	34	2	115
Annotator 1 (percentages)	10	22	62	4	0	2	100
Annotator 2 (percentages)	16.92	4.62	24.62	0	52.31	1.54	100
Both (percentages)	13.91	12.17	40.87	1.74	29.57	1.74	100
Agreement YES	5	2	12	0	0	1	20
Agreement NO	38	37	14	47	15	48	199
Disagreement YES	6	10	23	2	34	0	75
YES - Annotator 1 only	0	9	19	2	0	0	30
YES - Annotator 2 only	6	1	4	0	34	0	45
Percentage of agreement	87.76	79.59	53.06	95.92	30.61	100	74.49
Cohen's kappa	0.56	0.21	0.14	n/a	n/a	1	0.19

Table 3.7: Overall open card sorting results from two annotators (Annotator 1 and Annotator 2)

	Annotator 1	Annotator 2	Annotator 1 (percentages)	Annotator 2 (percentages)
#1	2	n/a	3.51	n/a
#2	13	4	22.81	3.45
#3	7	n/a	12.28	n/a
#4	5	5	8.77	4.31
#5	18	n/a	31.58	n/a
#6	12	n/a	21.05	n/a
#7	n/a	17	n/a	14.66
#8	n/a	3	n/a	2.59
#9	n/a	6	n/a	5.17
#10	n/a	3	n/a	2.59
#11	n/a	6	n/a	5.17
#12	n/a	2	n/a	1.72
#13	n/a	13	n/a	11.21
#14	n/a	9	n/a	7.76
#15	n/a	20	n/a	17.24
#16	n/a	2	n/a	1.72
#17	n/a	11	n/a	9.48
#18	n/a	4	n/a	3.45
#19	n/a	9	n/a	7.76
#20	n/a	2	n/a	1.72
Total	57	116	100	100

Table 3.8: Open group card sorting results from two annotators (Annotator 1 and Annotator 2), with cases of agreement and disagreement

	#2: Committee	#4: Market Exclusivity/Monopoly	Total
Annotator 1	13	5	18
Annotator 2	4	5	9
Both	17	10	27
%Annotator 1	22.81	8.77	31.58
%Annotator 2	3.45	4.31	7.76
%Both	9.83	5.78	15.61
Agreement YES	4	3	7
Agreement NO	36	42	n/a
Disagreement YES	9	4	n/a
YES - Annotator 1 only	9	2	11
YES - Annotator 2 only	0	2	2
Percentage of agreement	81.63	91.84	n/a
Cohen's kappa	0.40	0.55	n/a

Table 3.9: Network metrics of G1 and G2 graphs. #V = vertices, #E = edges, AvDegr = average degree, Dia = diameter, AvPaLe = average path length, Den = density, Tra = transitivity, AvClCo = average clustering coefficient.

Name	#V	#E	AvDegr	Dia	AvPaLe	Den	Tra	AvClCo
G1	49	320	13.1	3	1.729	0.272	0.394	0.645
G2	548	1,279	4.7	7	3.248	0.009	0.069	0.132

Table 3.10: Classification results for the existence of a relationship between two legislative segments in each feature set (Bigr200 and Trig200). LR = logistic regression, DT = decision tree, SVM = support vector machines, NB = naive Bayes, DU = dummy rule-based classifier.

Set	Performance	LR	DT	SVM	kNN	NB	DU
Bigr200	F-measure	0.838	0.792	0.838	0.752	0.488	0.656
	Precision	0.814	0.806	0.775	0.825	0.921	0.666
	Recall	0.869	0.781	0.914	0.695	0.334	0.709
	Accuracy	0.768	0.720	0.758	0.687	0.522	0.579
Trig200	F-measure	0.840	0.775	0.843	0.783	0.505	0.699
	Precision	0.806	0.783	0.777	0.819	0.949	0.680
	Recall	0.879	0.770	0.924	0.731	0.349	0.652
	Accuracy	0.768	0.696	0.765	0.701	0.538	0.57

Table 3.11: Classification results: F-measure values (average and standard deviation) concerning the existence of a relationship for each link norm type using Bigr200 features with different classification methods. LR = logistic regression, DT = decision tree, SVM = support vector machines.

Norm types	LR	DT	SVM
MO	0.753 (0.057)	0.715 (0.079)	0.775 (0.04)
II	0.789 (0.159)	0.760 (0.175)	0.550 (0.146)
VOL	0.753 (0.094)	0.755 (0.139)	0.686 (0.118)
PR	0.564 (0.124)	0.535 (0.143)	0.522 (0.168)

Table 3.12: Classification results on predicting the existence (or non-existence) of relationships between two legislative portions with only a bag-of-bigrams model (left column) and only network features (right column). LR = logistic regression, DT = decision tree, SVM = support vector machines, kNN = k-nearest neighbours.

Performance	Only bag-of-bigrams				Only network features			
	LR	DT	SVM	kNN	LR	DT	SVM	kNN
F-measure	0.828	0.828	0.832	0.726	0.808	0.777	0.808	0.757
Precision	0.809	0.821	0.776	0.840	0.698	0.798	0.691	0.71
Recall	0.852	0.839	0.900	0.650	0.962	0.760	0.979	0.819
Accuracy	0.756	0.763	0.751	0.670	0.689	0.701	0.684	0.644

Chapter 4

A lightweight ontology and a harmonised glossary for multilingual and analogical interpretation

This chapter presents two innovative methodologies leveraging knowledge modelling techniques to address the challenge of multilingualism in EU legal harmonisation and the proper alignment of EU legislative acts and corresponding national implementing legislation. Firstly, it introduces an analogical lightweight ontology designed to mitigate discrepancies arising from translations of EU legislative acts across the 24 official EU languages, thus promoting comparative and analogical interpretation. Secondly, the chapter details the creation of a harmonised glossary within the EU-funded Facilex project. This glossary significantly enhances the ability of LLMs to automatically identify, align and compare domestic legislative concepts with their EU counterparts, thereby facilitating more accurate and efficient analysis of legal harmonisation efforts. The integration of these tools

improves the semi-automation of concept extraction and legislative alignment, pushing the boundaries of current state-of-the-art techniques in legal technology.

4.1 Related work

Ontologies have been widely adopted within legal informatics to structure and interpret legal knowledge for both human and automated processes [280]. Since the 2010s, numerous ontologies have been developed targeting specific legal sub-fields, each designed to have specific characteristics and tools that make them uniquely adaptable to particular objectives.

Instead of merely transcribing normative elements like permissions, prohibitions, and duties, the ontology presented in this chapter defines concepts pertaining to criminal procedural rights within the European Union. This approach leverages the Semantic Web, acknowledged as the premier means for managing knowledge through standardised formats such as the Resource Description Framework (RDF) and the Web Ontology Language (OWL), which are readily accessible online [281].

The Open Digital Rights Language (ODRL) [282], developed by the ODRL Community Group [283], delineates digital content and media policies [284]. This framework consists of a Core Vocabulary designed to structure policies and a Common Vocabulary that outlines general terms to define actions including obligations, permissions, and prohibitions.

Expanding on ODRL, the Linked Data Rights [285] (LDR) ontology, formulated by the Ontology Engineering Group [286], enhances ODRL's foundational elements such as Action, Asset, Policy, and Rule to govern the usage conditions of Linked Data.

In related legal contexts, the Creative Commons Rights Expression Language (ccREL) [287] provides a standard for articulating copyright licensing terms in a format interpretable by machines. Similarly, the Licence

for Linked Open Data [288] (L4LOD) employs a streamlined ontological framework to manage licensing terms within the Data Web. Here, the deontic categories of permission, prohibition, and obligation are distinctly marked to guide actions permitted and prohibited within licensing terms.

Focusing on public procurement, the LOTED2 ontology [289] developed by Distinto et al. [290] aims to encapsulate information pertinent to public procurement within the European Union, utilising terminology from Tenders Electronic Daily (TED) [291], a database that aggregates procurement notices from the EU and member states public institutions.

Conversely, Muñoz-Soro et al.'s [292] Public Procurement Ontology [293] (PPROC) semantically depicts the details found in Spanish and EU official legal procurement documents. PPROC aims to chart the progression of tenders from announcement to completion and, among other features, includes a taxonomy of contracts utilised in procurement processes.

Many ontologies are designed to address the data protection and privacy law sectors. Pandit et al.'s [294] GDPRtEXT [295] (GDPR text extensions) encapsulates General Data Protection Regulation No. 2016/679 (GDPR) by representing it as a Linked Data resource. Each significant segment of the GDPR is associated with a unique Uniform Resource Identifier (URI) to ensure comprehensive visualisation of the regulation's framework. This includes the depiction of articles, recitals, citations, and the interconnections within its provisions.

The elements of the GDPR are further explored in Palmirani et al.'s [296] Privacy Ontology (PrOnto). PrOnto not only facilitates information retrieval but is also grounded in a robust theoretical model that encompasses legal reasoning and compliance verification techniques.

Conversely, Oltramari et al.'s [297] (2018) Privacy Ontology (PrivOnto), developed as part of the Usable Privacy Policy Project [298], focuses on modelling annotated privacy policies that detail the data protection practices of websites. In summary, numerous endeavours have successfully established cross-domain ontologies to counteract the fragmentation ob-

served in legal domains.

EuroVoc [299], maintained by the Publications Office of the European Union, serves as a multilingual, multidisciplinary thesaurus that indexes legal and political documents from EU institutions, enhancing their accessibility.

While not strictly an ontology, LegalRuleML [300] by Palmirani et al. [301] and Athan et al. [302] provides a framework for the representation and interchange of legal knowledge. LegalRuleML facilitates the synchronisation of diverse legal documents such as laws and official policy documents. Similarly, the European Legislation Identifier [303] (ELI) standardises the publication of legal texts with a uniform metadata set, promoting interoperability among national administrations.

Building upon LegalRuleML, the Normative Requirements Vocabulary [304] (NRV) by Gandon et al. [305] utilises Semantic Web standards to model normative rules and requirements.

Besides ontological modelling, significant research has been devoted to automating semantic relationship extraction and taxonomy creation, focusing on defining modelling and automated detection tasks. Nevertheless, most studies in NLP and ML prioritise the identification of lexico-syntactic patterns with neural architectures [306, 307, 308], whereas more intricate analogical reasoning requires novel investigative approaches and future exploration.

Moreover, while current language technologies that concentrate on content extraction, detection and labelling face limitations in reasoning with and handling domain-specific information, semantic-labelled knowledge tends to enhance these technologies. The integration of semantic knowledge remains a challenging issue that warrants further investigation [309].

Meanwhile, the recent surge in natural language processing (NLP) performance is intrinsically linked to the development of the transformer architecture [310]. This foundational work enabled subsequent models like bidirectional encoder representations from transformers (BERT) [311] and

ultimately led to the current generation of large language models (LLMs). Models such as Generative Pre-trained Transformer (GPT) [312], Large Language Model Meta AI (Llama) [313], Gemini [314], and Claude [315] now exhibit impressive proficiency across a range of complex language tasks [316]. The application of these advancements is demonstrably impacting legal practice, academic research [317, 318] and commercial applications [319], signifying a transformative approach to complex language data analysis and deployment across diverse domains.

At the same time, investigations into the evaluation of LLM-based methodologies have compared manual and automated techniques. Recent comparative analyses underscore the critical importance of human evaluation, particularly for tasks demanding expertise in specialist subject-matters [320]. Most interestingly, LEGALBENCH is a collaboratively built benchmark for evaluating legal reasoning in LLMs. It consists of 162 tasks across six categories of legal reasoning: issue-spotting, rule-recall, rule-application, rule-conclusion, interpretation, and rhetorical-understanding. The benchmark was created by legal professionals to ensure practical relevance and for evaluating 20 LLMs [321].

Future studies could benefit from an integrated analysis that combines these diverse perspectives to tackle the complexities of legal harmonisation in an evolving legal, cultural, and technological landscape.

4.2 An analogical lightweight ontology for EU criminal procedural rights in judicial cooperation

This section presents the development of a streamlined ontology dedicated to criminal procedural rights within the realm of judicial cooperation. The ontology is aimed at assisting legal professionals by elucidating the specific contextual meanings of terms to facilitate the creation of a rule-based

ontology to model criminal procedural rights in judicial contexts.

Determining the applicable scope of legislative provisions presents significant challenges. Although not all legislative acts provide explicit definitions, usually enshrined in the initial articles, experts such as judges engage in intricate legal reasoning to elucidate and justify the application of norms in various cases.

In pursuit of this goal, judicial bodies employ interpretative methodologies. Among these, systematic interpretation is prominent, which seeks to ascertain the meanings of legal norms “by examining the law solely within its legislative (Gesetz, lois) or other pertinent legal frameworks, meaning within one or several legal texts belonging to the same legal system” [322].

Systematic interpretation typically relies on analysing norms from the legal text or referring to supplementary sources. For instance, the Court of Justice of the European Union (CJEU) frequently bases its reasoning on the recitals that precede the normative provisions in EU laws. Consequently, when legal texts lack specific articles that define concepts, their meanings are often inferred from other legislative texts that include the relevant information.

Nevertheless, some legislative acts contain “implicit” definitions with connecting keywords such as “for instance”. Moreover, definitions often encompass connecting keywords such as “for instance”, “for example”, “including”, “excluding”, “such as”, which makes interpretation challenging. To address this, an “analogical” ontology has been developed to aid legal professionals in analogical reasoning, especially when such definitions provide tangible examples of the concepts defined. This objective is facilitated through ontology engineering, a mature field within legal informatics that enhances the representation and comprehensibility of legal texts.

The granularity of ontologies can differ significantly. For example, the Legal Knowledge Interchange Format (LKIF) Core Legal Ontology [323] is designed to be jurisdiction agnostic, the Lexical Ontologies for Legal Information Sharing (LOIS) ontology framework [324] distinguishes between

EU terms and national legal terms, and the European Legal Taxonomy Syllabus (ELTS) ontology framework [325] employs a bottom-up approach that supports multiple term definitions in each jurisdiction, each linked explicitly to the source of that definition. Definitions in the ELTS ontology typically follow a structured format such as “*X means Y*”, “*X is defined as Y*”, or “*X pertains to Y*”.

However, there is a notable gap in ontological studies concerning the representation of “implicit” definitions, which sometimes consist of dispersed segments within statutory clauses. The objective is to aggregate these conceptual fragments within laws to outline comprehensive definitions that aid in judicial decision-making, particularly when judges need to rely on multiple legal foundations to substantiate their rulings. Thus, the research described in this chapter contributes to the advancement of legal ontologies by extending the types of definitions discussed in previous studies [326], showing how these can be applied in the manual identification of implicit definitions, and developing an annotated dataset for future research, including automated categorisation of paragraphs, and a repository of NLP methods for each type of definition.

A comprehensive review on the nature of definitions [327] indicated that traditional definitions often include hypernyms (typically broad rather than specific), meronyms, synonyms, and purpose-oriented details. Within the *CrossJustice* initiative¹, a distinct challenge arose from the fact that among the six relevant directives, only two featured a dedicated article for definitions. Specifically, Article 3 of Directive 2016/800 provides three definitions: for the terms “*child*”, “*holder of parental responsibility*”, and “*parental responsibility*”. Meanwhile, Article 3 of Directive 2016/1919 only defines the term “*legal aid*”. Although classical definitions exist elsewhere within these directives and are utilised in the ontology, they are sparse. For example, recital 15 of Directive 2013/48/EU describes the term “*lawyer*” as follows: “*The term ‘lawyer’ in this Directive refers to any person who, in accordance with national law, is qualified and entitled, including by means*

¹Grant Agreement n.847346 - <https://www.crossjustice.eu/en/index.html>

of accreditation by an authorised body, to provide legal advice and assistance to suspects or accused persons.” These definitions, often absent from conventional sections, also employ different connective keywords to those typically used in classical definitions.

This situation presented an unique opportunity to investigate a commonly overlooked aspect in the creation of domain-specific legal ontologies. Regardless of the presence or absence of classical definitions, legal texts often contain many unmarked or incomplete definitions that could still aid legal professionals (and automated legal reasoning systems) in drawing analogical or teleological conclusions. Such definitions are not limited to articles but are also found in recitals, which are crucial for the legal interpretations made by the Court of Justice of the European Union (CJEU).

In [328], various types of these definitions were identified and categorised as follows:

- *Example Definitions*: these describe a concept through typical examples, encouraging analogical reasoning. They convey a sense of completeness, suggesting that applicable instances must align closely with the examples provided or their analogues.
- *Include or Exclude Definitions*: these definitions typically clarify the inclusion or exclusion of certain examples of a particular concept in situations where their inclusion nor exclusion might otherwise be ambiguous or unexpected. These definitions are inherently incomplete, as additional examples may or may not fall within the scope of the term under consideration.
- *Definitions by Reference*: certain laws explicitly reference other legislation to obtain their definitions of specific terms. In this way, these external definitions also apply to the referring documents, thereby broadening the scope of the previous legislation.

In developing the *CrossJustice* ontology, these definition classes were refined and additional classes were introduced.

Applying the law inherently involves applying abstract principles to particular situations. A terminology-centric legal ontology serves as a valuable resource for understanding the meaning of terms and their relationships, enhancing search capabilities, and structuring rule-based ontologies. Legal practitioners could significantly benefit from a tool that consolidates the essential attributes of legal concepts, offers precise definitions, and enhances the predictability of judicial decisions.

4.2.1 Tools and methodology

The ontology discussed in this chapter covers definitions associated with terms derived from six EU directives:

- Directive 2010/64/EU of the European Parliament and of the Council, dated 20 October 2010, which addresses the right to interpretation and translation during criminal proceedings;
- Directive 2012/13/EU of the European Parliament and of the Council, enacted on 22 May 2012, focusing on the right to be informed in criminal proceedings;
- Directive 2013/48/EU of the European Parliament and of the Council, issued on 22 October 2013, which ensures the right to legal counsel in criminal proceedings and European arrest warrant proceedings, the right to notify a third party upon arrest, and the right to communicate with third parties and consular officials while in custody;
- Directive (EU) 2016/343 of the European Parliament and of the Council, from 9 March 2016, which enhances certain aspects of the presumption of innocence and the right to participate in one's own trial;
- Directive (EU) 2016/800 of the European Parliament and of the Council, dated 11 May 2016, concerning procedural protections for children who are suspects or accused in criminal proceedings;

- Directive (EU) 2016/1919 of the European Parliament and of the Council, issued on 26 October 2016, pertaining to legal aid for suspects and accused individuals in criminal proceedings, as well as for persons subject to European arrest warrant proceedings.

In addition to the definitions from these six directives, the dataset compiled for this ontology also incorporates relevant segments from other legal sources, including EU and international treaties, EU legislative acts, policies, cited within the directives and all relevant judgments from the Court of Justice of the European Union that offer critical interpretations of the clauses of the directives.

Theoretical basis

In an effort to compile and depict definitions within a minimalistic ontology, the classification system from [328] was initially followed. These definition types were further refined and additional types were introduced. Consequently, we have adopted the following classes of definitions:

- The **Classical (or Regular) Definition**² represents the conventional form of definitions, typically featuring formulaic phrases such as “*X means Y*” or “*X is understood to mean Y*”. Di Caro [327] notes that such definitions usually include synonyms, hypernyms, meronyms, and purpose-related information. The hallmark of Classical Definitions is their comprehensive nature. For instance, Article 1(1) of Directive 2010/64 declares, “*This Directive establishes rules concerning the right to interpretation and translation in criminal proceedings and in the context of European arrest warrant proceedings.*” From this, the Classical Definition of Directive 2010/64 was derived as: “EU regulation establishing guidelines for the right to interpretation and translation in criminal and European arrest war-

²The term SenseDefinition from the Linked Term Bank of Copyright-Related Terms was set aside in favour of the term Classical Definition, as used in [328].

rant proceedings.” Such definitions are characterised by their clarity, depth, and readability [327].

- The **Part Definition** encapsulates the components or elements of a concept, such as procedures or rights, where the collective meaning emerges from these components. An illustrative example is found in Article 4(2) of Directive 2012/13, where each listed item forms a part of the set of rights about which information must be provided in a Letter of Rights. Beyond the rights outlined in Article 3, the Letter of Rights should include information on the following, as applicable per national law:

1. the right of access to the materials of the case;
2. the right to have consular authorities and one person informed;
3. the right of access to urgent medical assistance; and
4. the maximum number of hours or days suspects or accused persons may be deprived of liberty before being brought before a judicial authority.

These elements, detailed within various normative clauses of Directive 2012/13, collectively define the Letter of Rights and serve as a reference point for relevant parties.

- The **Essential Part Definition** identifies critical components or elements necessary for the concept’s existence. For example, recital 33 of Directive 2016/800 states: “*The confidentiality of communications between children and their lawyer is fundamental to the effective exercise of defence rights and is an integral part of the right to a fair trial.*” The phrases “*is fundamental to*” and “*is an integral part of*” suggest this is an instance of EssentialPartDefinition.
- The **Purpose Definition** elucidates a concept through its intended purpose. For instance, in Article 7(4) of Directive 2012/13, denial of access to certain materials may be permitted under specific conditions: “*By way of derogation from paragraphs 2 and 3, provided*

that this does not prejudice the right to a fair trial, access to certain materials may be refused if such access may lead to a serious threat to the life or the fundamental rights of another person or if such refusal is strictly necessary to safeguard an important public interest, such as in cases where access could prejudice an ongoing investigation or seriously harm the national security of the member state in which the criminal proceedings are instituted.” Accordingly, the more specific purposes of 1) preventing interference with ongoing investigations and 2) protecting the national security of the member state conducting the criminal proceedings were also recognised.

- The **Parameter Definition** involves specific factors considered in the application of a legal concept, thus enhancing comprehension of that concept. A prime example from Article 8(2) of Directive 2016/800 illustrates a Parameter Definition: *“The results of the medical examination shall be taken into account when determining the capacity of the child to be subject to questioning, other investigative or evidence-gathering acts, or any measures taken or envisaged against the child”*.
- The **Ratione Temporis Definition** defines the temporal scope for which a legal definition, such as a principle, right, or obligation, is applicable. Article 2(1) of Directive 2016/800 captures a Ratione Temporis Definition: *“This Directive applies to children who are suspects or accused persons in criminal proceedings. It applies until the final determination of the question whether the suspect or accused person has committed a criminal offence, including, where applicable, sentencing and the resolution of any appeal”*.
- The **Ratione Personae Definition** specifies the legal subjects affected by a principle, right, or obligation. Article 2 of Directive 2016/343 articulates this by stating: *“This Directive applies to natural persons who are suspects or accused persons in criminal proceedings. It applies at all stages of the criminal proceedings, from the*

moment when a person is suspected or accused of having committed a criminal offence, or an alleged criminal offence, until the decision on the final determination of whether that person has committed the criminal offence concerned has become definitive”.

- The **Typical Example Definition** (a subclass of the Example Definition) uses a common instance of a broader concept to define it. For example, Article 2(3) of Directive 2010/64 states: “[t]he right to interpretation under paragraphs 1 and 2 includes appropriate assistance for persons with hearing or speech impediments.”.
- The **Atypical Example Definition** (another subclass of Example Definitions) refers to an example that is unusually part of a broader concept. For example, recital 16 of Directive 2016/800 states: “*is understood to mean the final determination of the question whether they have committed the offence, including, where applicable, sentencing and the resolution of any appeal*”. This clarification that the proceedings’ conclusion includes appeal stages, not typically considered part of the process, represents an atypical example. The use of “include” as a connecting keyword can suggest either a Typical or Atypical Definition depending on the context.
- The **Important Example Definition** (also a subclass of Example Definitions) highlights an essential example of a broader concept and underscores that its inclusion is mandatory. For instance, recital 13 of Directive 2013/48 stresses the care due to suspects or accused persons, especially “*towards individuals with any physical impairments affecting their communication abilities*”.
- The **Parameter Example Definition**, subsuming characteristics of both *ExampleDefinition* and *ParameterDefinition*, uses examples of parameters for conceptual clarification. However, like other *ExampleDefinitions* discussed, its parameter set is incomplete, thus encouraging reasoning by analogy. This approach is mirrored in Recital 4 of Directive 2013/48/EU. “*Under the Treaty on the European Union*

(TEU), specifically Article 5, the EU may implement actions adhering to the subsidiarity principle. This Directive is crafted following the proportionality principle outlined in the same article, ensuring that it does not exceed what is necessary to fulfill its objectives.”

- The **Non Example Definition** (a subclass of Example Definitions) utilises examples that are specifically excluded from a broader category to sharpen the definition of that category. An illustration of this is found in recital 13 of Directive 2013/48, which clarifies what does not constitute “*criminal proceedings*” by excluding certain types of proceedings, thus refining the concept. It states: “*proceedings in relation to minor offending which take place within a prison and proceedings in relation to offences committed in a military context which are dealt with by a commanding officer should not be considered to be criminal proceedings for the purposes of this Directive*”.
- The **Definition By Reference** highlights that not all laws provide definitions for every term; sometimes, they explicitly direct readers to definitions in other legislative documents. For instance, recital 49 of Directive 2016/343 explains, “*the Union may adopt measures in accordance with the principle of subsidiarity as set out in Article 5 TEU [Treaty on the European Union]. In accordance with the principle of proportionality, as set out in that article, this Directive does not go beyond what is necessary in order to achieve those objectives.*” The TEU definitions of the subsidiarity and proportionality principles are thereby explicitly applied to Directive 2016/343.

As documented in Table [4.2](#), beyond Classical (or Regular) definitions, the taxonomy includes other definition classes categorised under Analogical Definitions. Generally, these Analogical Definitions employ examples as concrete manifestations of abstract concepts to help undersand what they mean.

The employment of example-based definitions is apparent due to their explicit usage of illustrative examples. The Analogical category was ex-

tended to encompass other classes based on the annotation outcomes, which suggest that the Part, Essential Part, Purpose, Parameter, Ratione Temporis, and Ratione Personae classes share lexical features conducive to enhancing analogical reasoning and facilitating thorough interpretations in legal analysis, particularly judicial interpretations. Analogical Definitions, especially when combined with NLP tools, are critical for advancing understanding and application of legal concepts [329].

This study is also guided by principles from the European Legal Taxonomy Syllabus (ELTS) [325], which influences the approach as follows:

- Definitions are presumed to pertain to the legislative source itself unless specifically stated otherwise. Scope restrictions were recognised with phrases like “*for the purposes of paragraph X*” and scope expansions were identified through explicit references to definitions from other legislation.
- It is presumed that definitions are jurisdiction-specific. In the EU context, the transposition of legislation—and the concepts defined therein—might alter the definitions, necessitating clear delineation of the relationship between associated concepts.

Ontological framework

The construction of the ontology leveraged the Linked Term Bank of Copyright-Related Terms [330]. This ontology, which is tailored to a specific legal domain distinct from that of the ELTS, is both multilingual and multi-jurisdictional. The intuitive organisation of its concepts and terms, which follows established best practices in ontology development, is well-appreciated. The foundation of the Copyright Term Bank is based on LExicon Model for ONtologies (lemon) and Simple Knowledge Organization System (SKOS) classes. The Linked Term Bank of Copyright-Related Terms [330] were integrated into WebProtégé [331] to facilitate an analysis of its structural components and to inform the development process.

In developing the lightweight ontology, the classes from the Copyright Term Bank were adopted, as detailed in Table 4.1.

Class	Description
Concept	the definiens
LexicalEntry	the words or phrases used to represent the Concept.
LexicalSense	the lexical meaning of a lexical entry. When linked to a Concept, this implies that the lexical entry can be used to refer to that Concept.
ClassicalDefinition	the definiendum, along with the legal source of the definiendum.

Table 4.1: Classes inherited from the Copyright Term Bank

To enhance the framework adopted for the ontology, the classes outlined in Section 3.1 have been incorporated. The types of definitions that have been newly identified in this research are enumerated in Table 4.2).

Type of class	<i>Class</i>
Analogical	<i>Part</i>
Analogical	<i>EssentialPart</i>
Analogical	<i>Purpose</i>
Analogical	<i>Parameter</i>
Analogical	<i>RationeTemporis</i>
Analogical	<i>RationePersone</i>
Analogical	<i>ParameterExample</i>
Analogical	<i>TypicalExample</i>
Analogical	<i>AtypicalExample</i>
Analogical	<i>ImportantExample</i>
Analogical	<i>Non-Example</i>

Table 4.2: New analogical classes created for the analogical lightweight ontology

Given that the newly identified definition types inherently entail relations between concepts, these relations have been modelled and are detailed in Table 4.3.

The duplication offers these benefits:

Is-Relationships	Has-Relationships
IsPartOf	HasPart
IsPurposeOf	HasPurpose
IsParameterOf	HasParameter
IsRationeTemporis	HasRationePersone
IsTypicalExampleOf	HasTypicalExample
IsImportantExampleOf	HasImportantExample
IsParameterExample	HasParameterExample
IsNonExampleOf	HasNonExample
IsEssentialPartOf	HasEssentialPart

Table 4.3: Is- and Has-Relationships related to the definition types in the analogical lightweight ontology

1. it facilitates direct access to the original source text within the defined instances;
2. it allows users to clearly see how various concepts are interconnected, with reference to the pertinent legal source.

4.2.2 Data collection and analysis

A legal expert examined six European directives concerning criminal procedural rights in judicial cooperation to identify “classical” and “analogical” or “non-classical” definitions as described above. This expert also reviewed rulings from the Court of Justice of the European Union to extract additional definitions that would enrich the ontology. Further, the expert included definitions from EU treaties and charters such as the TEU, TFEU, CFREU, and from international conventions like the European Convention on Human Rights (ECHR), the International Covenant on Civil and Political Rights (ICCPR), and the Vienna Convention on Consular Relations, all of which are referenced directly in the directives and categorised as DefinitionByReference in the accompanying Excel table.

The legal expert initially examined the directives in English, then compared the relevant legal texts with their counterparts in Italian, French, and German. While most provisions showed little variation in meaning

across languages, a few discrepancies were noted.

For example, the English text of Article 2(1) of Directive 2010/64 reads: *“member states shall ensure that suspected or accused persons who do not speak or understand the language of the criminal proceedings concerned are provided, without delay, with interpretation during criminal proceedings before investigative and judicial authorities, including during police questioning, all court hearings and any necessary interim hearings.”*

Conversely, the Italian text states: *“Gli Stati membri assicurano che gli indagati o gli imputati che non parlano o non comprendono la lingua del procedimento penale in questione siano assistiti senza indugio da un interprete nei procedimenti penali dinanzi alle autorità inquirenti e giudiziarie, inclusi gli interrogatori di polizia, e in tutte le udienze, comprese le necessarie udienze preliminari.”*

This delineates a subtle distinction in how the concepts are linked. In the English version, the expression “criminal proceedings before investigative and judicial authorities” includes three scenarios: “police questioning”, “all court hearings”, and “any necessary interim hearings”. In contrast, the Italian version frames the third scenario as a particular instance of the second. The French and German versions, however, mirror the structural composition observed in the English text.

Comparing linguistic versions can also reveal uncertainties and ambiguities. For example, recital 28 of Directive 2010/64 states: *“[w]hen using videoconferencing for the purpose of remote interpretation, the competent authorities should be able to rely on the tools that are being developed in the context of European e-Justice (e.g. information on courts with videoconferencing equipment or manuals).”* It may seem odd to categorise “information” as a type of “tool”. Nevertheless, when the different language versions are examined, the intended meaning becomes evident.

For instance, the Italian version of the recital reads: *“Quando si utilizza la videoconferenza per l’interpretazione a distanza, le autorità competenti dovrebbero poter utilizzare gli strumenti sviluppati nel contesto della*

giustizia elettronica europea (ad esempio informazioni sui tribunali che dispongono di materiale o di manuali per la videoconferenza)."

Linguistic analysis has also revealed that recital 28 cites only one example of tools, "*information on courts with videoconferencing equipment or manuals*", instead of listing two separate instances: "*information on courts with videoconferencing equipment*" and "*manuals*" — a distinction that could be misinterpreted in the English version alone.

The German translation of Article 6(3) of the ECHR explicitly highlights differences between the English and French texts: "*Jeder Angeklagte hat **mindestens (englischer Text)** [emphasis added] **insbesondere (französischer Text)** [emphasis added] die folgenden Rechte...*". Notably, while the English version of Article 6(3) articulates that "[e]veryone charged with a criminal offence has the following minimum rights: [...]", the French translation substitutes "*minimum*" with "*notamment*", translating to "in particular". The German text uses "*mindestens*" to reflect "*minimum*" from the English version and "*insbesondere*" as the equivalent of "*notamment*", which also means "*in particular*". The Italian version employs "*in particolare*", which also means "*in particular*". The inclusion of these variations in the German text hints at a potential semantic divergence between "*minimum*" and "*in particular*". While these differences are presumably more lexical than semantic, further investigation is warranted to determine if these discrepancies influence diverse domestic implementations of the right to a fair trial under Article 6 of the ECHR. However, this multilingual analysis is still in progress, with the potential for new insights in future work.

In the initial phase of this research, the legal expert annotated articles and recitals that featured typical connective keywords indicative of various definition types, such as "*include*", "*such as*", "*inter alia*", and "*in particular*". This selection intentionally included negative examples — provisions with these keywords that did not meet the criteria for definitions as previously specified — to assist in training a machine learning classifier in automated ontology population. This keyword search proved

mostly beneficial for identifying Example Definitions.

The next phase required more in-depth analysis; the complete text of the directive was scrutinised for any instances of definitions. Identifying classical definitions was particularly challenging with the directives under consideration as they were rarely in the conventional format of European directives and there were no articles specifically dedicated to classical definitions. The findings suggest that a semi-automated approach may be more effective for some types of definitions than others.

In a subsequent phase, a second annotator with expertise in legal informatics conducted a review to ensure no definitions were overlooked. This step occurred after the initial concept extraction process (detailed below) as the understanding of definition types, particularly Purpose Definitions and Parameter Definitions, evolved and deepened. This process, reflective of the innovative and iterative nature of the work described in this chapter, necessitated continual refinement.

Table 4.4 and Figure 4.1 provide numerical and visual representations of the outcomes from the definition extraction and verification phases.




Figure 4.1: Distribution of definition types over the six directives. SenseDefinition stands for ClassicalDefinition for the purpose of this research work.

In this research work, definitions from both recitals and articles were

Def. Type	<i>D.64</i>	<i>D.13</i>	<i>D.48</i>	<i>D.343</i>	<i>D.800</i>	<i>D.1919</i>	Total
<i>Classical</i>	3	7	19	8	15	4	56
<i>By reference</i>	12	17	15	15	3	5	67
<i>Typical Ex.</i>	5	4	24	13	19	5	70
<i>Atypical Ex.</i>	2	0	3	0	7	1	13
<i>Import. Eex.</i>	2	2	6	1	10	2	23
<i>Non-example</i>	1	0	2	10	5	1	19
<i>Part</i>	1	0	6	3	11	2	23
<i>Essential part</i>	2	2	2	2	2	0	10
<i>Purpose</i>	4	3	3	3	4	1	18
<i>Rat. temp.</i>	1	1	7	2	2	0	13
<i>Rat. pers.</i>	1	0	1	2	1	1	6
<i>Parameter</i>	0	0	0	0	3	1	4
<i>Parameter Ex.</i>	1	1	2	1	0	0	5
Total	35	37	90	60	82	23	327

Table 4.4: Number of definitions in the six directives that were extracted and modelled. D.64 = Directive 2010/64/EU, D.13 = Directive 2012/13/EU, D.48 = Directive (EU) 2013/48/EU, D.343 = Directive (EU) 2016/343, D.800 = Directive (EU) 2016/800, D.1919 = Directive (EU) 2016/1919. Typical Ex. = Typical Example, Atypical Ex. = Atypical Example, Rat. temp. = Ratione Temporis, Rat. personae = Ratione Personae, Parameter Ex. = Parameter Example. Parameter Example

included in the corpus of instances. From the outset, the greater number of recitals compared to articles was noted, as well as their greater length (Table 4.5).

In the ontology, the source of each definition is explicitly cited, allowing users to assess the significance of each source according to their own judgment. During discussions, there were instances where it was difficult to determine how to represent certain normative provisions. A straightforward colour-coding system was adopted to track progress: green indicated “agreed”, orange indicated “under discussion”, and red was used for normative agreements that lacked definitions.

Owing to other commitments, it was not feasible to continue working collaboratively on all the directives. Consequently, in a later phase, the legal informatics annotator extracted several concepts, and her work was subsequently reviewed by the legal expert annotator.

Directive	Total number of recitals per directive	Total number of articles per directive	Total word count of recitals per directive (English version)	Total word count of articles per directive (English version)
2010/64	36	12	2314	1525
2012/13	45	14	3118	1574
2013/48	59	18	6030	2721
2016/343	51	16	4193	1565
2016/800	71	27	6269	4683
2016/1919	33	14	2737	1515

Table 4.5: Number of articles and recitals per directive and total word count of recitals and articles

4.2.3 Ontology creation

The WebProtégé software [331] was selected for creating the ontology due to not only its renowned capabilities in ontology creation but also its collaborative features. In practice, while the tool proved relatively user-friendly, difficulties were encountered in modifying entries once they were made. Specifically, it was challenging to correct or delete fields that contained errors (see Figure 4.2).




Figure 4.2: Example of an error that cannot be easily corrected in WebProtégé

Due to these limitations, all the preliminary analysis was conducted

using Excel and, the WebProtégé ontology was only populated as a final step. The first step was detailing the structure of the ontology, including specific examples.

Initially, the Linked Term Bank of Copyright-Related Terms [330] was imported into WebProtégé. This term bank was available as an N-Triples file [332], which we converted to the RDF format using the EasyRDF online tool [333]. After importing the RDF file into WebProtégé, we faced issues with visibility: none of the ClassicalDefinition instances were visible, and the relationships between concepts and their ClassicalDefinitions had not been maintained. Nevertheless, the primary objective was not to reuse the data but to understand the general structure, which we mainly achieved using the WebProtégé tool, supplemented by the N-Triples and RDF files as necessary. The structure of the Linked Term Bank of Copyright-Related Terms is summarised as follows:

- *Owl:Thing* includes four direct subclasses: *Concept*, *LexicalEntry*, *LexicalSense*, and *SenseDefinition*;
- Concepts are associated with several *AnnotationProperties*, including:
 - *rdfs:label*: the commonly used name for this *Concept* as a *plain-Literal* value;
 - *skos:definition*: links to an instance of *SenseDefinition*, which includes the definition, source, and other pertinent details, and which we refer to as ClassicalDefinition;
 - *isSenseOf*: links to one or more *LexicalEntry* instances that articulate the *Concept*;
 - *jurisdiction*: links to a DBpedia entry detailing the jurisdiction;
 - *reference*: links to a DPBedia entry containing information about the jurisdiction;
 - *closeMatch*: connects to a similar concept in the InterActive Terminology for Europe (IATE) EU terminology database;

- *narrower*: links to narrower *Concept* instances;
- *rdfs:comment*: as a *plainLiteral*.

The *AnnotationProperties* *rdfs:label*, *skos:definition*, and *isSenseOf* are present in all *Concepts*.

- *LexicalEntries* include the following *AnnotationProperties*:
 - *rdfs:label*: a term describing a *Concept* in *plainLiteral*;
 - *denotes*: links to one or more *Concept* instances denoted by the term in the *LexicalEntry*;
 - *language*: the language of the term, represented as a *plainLiteral*;
 - *sense*: an *owl:NamedIndividual* of class *LexicalSense*).

The properties *rdfs:label*, *skos:denotes*, and *sense* are consistent across all *LexicalEntries*.

- *LexicalSenses* include the *AnnotationProperty*:
 - *reference*: linking to one or more instances of the class *Concept*;
- *SenseDefinitions*, renamed as *ClassicalDefinitions* in the analogical ontology, possess properties such as:
 - *source*: the name of the glossary from which the definition originates, including a URI for the glossary;
 - *value*: a *plainLiteral* value of the definition, with a language specifier.

In developing the ontology, all previously mentioned classes and properties were retained (with *SenseDefinitions* renamed as *ClassicalDefinitions*) because the main objective includes depicting classical definitions. In addition to *ClassicalDefinitions*, the ontology was expanded to include various other definition types, organising the class structure as follows:

- *Concept*
- *LexicalSense*
- *LexicalEntry*
- *Definition*
 - *ClassicalDefinition*
 - * *PartDefinition*
 - * *EssentialPartDefinition*
 - * *PurposeDefinition*
 - * *ParameterDefinition*
 - * *RationeTemporisDefinition*
 - * *RationePersonaeDefinition*
 - *AnalogicalDefinition*
 - * *TypicalExampleDefinition*
 - * *AtypicalExampleDefinition*
 - * *ImportantExampleDefinition*
 - * *ParameterExampleDefinition*
 - * *NonExampleDefinition*

Here is an instance of a Typical Example definition taken from Article 2(3) of Directive 2010/64:

The right to interpretation under paragraphs 1 and 2 includes appropriate assistance for persons with hearing or speech impediments.

In the ontology, the Concept “*the right to interpretation*” is linked to a *TypicalExampleDefinition*, which includes a definition field as well as a

Individual: the right to interpretation_EU

IRI

http://webprotege.stanford.edu/the_right_to_interpretation_EU

Annotations

rdfs:label

the right to interpretation_EU

isSenseOf

the right to interpretation-en

Enter property

Enter value

Types

☒ Concept

Enter a class name

Relationships

☐ HasExampleDefinitic

the right to interpretation_EU-exampledefinition

☐ HasTypicalExample

appropriate assistance for persons with hearing or speech impediments_EU

Enter property

Enter value

Same As

Enter an individual name

Figure 4.3: Representation of the “*right to interpretation*” in the analogical lightweight ontology

comment field that cites the original article for reference. Another Concept exists for “*appropriate assistance for persons with hearing or speech impediments*”.

The Copyright TermBank predominantly utilises annotation properties to show connections between concepts, their lexical senses, lexical entries, and sense definitions. The ontology incorporates definitions characterised by their interrelationships with other Concepts. Relationship properties were utilised to articulate these connections, enhancing the ability to visualise how different concepts are related to one another, as depicted in Figure 4.4.

Below is an *ImportantExampleDefinition* from recital 27, EU Directive 2010/64:

The duty of care towards suspected or accused persons who are in a potentially weak position, in particular because of any physical impairments which affect their ability to communicate effectively, underpins a fair administration of justice. The prosecution, law enforcement and judicial authorities should therefore ensure that such persons are able to exercise effectively the




Figure 4.4: Illustration of Relationship Properties for an *ImportantExampleDefinition*.

rights provided for in this Directive, for example by taking into account any potential vulnerability that affects their ability to follow the proceedings and to make themselves understood, and by taking appropriate steps to ensure those rights are guaranteed.

IRI
http://webprotege.stanford.edu/suspected_or_accused_persons_who_are_in_a_potentially_weak_position_EU

Annotations

- rdfs:label suspected or accused persons who are in a potentially weak position_EU
- isSenseOf suspected or accused persons who are in a potentially weak position-en

Enter property Enter value

Types

- ☒ Concept

Enter a class name

Relationships

- ☒ HasExampleDefinitic suspected or accused persons who are in a potentially weak position_EU-exampledefinition
- ☒ HasImportantExamp suspected or accused persons who have any physical impairments which affect their ability to communicate effectively_EU

Enter property Enter value

Figure 4.5: Depiction of “*suspected or accused persons who are in a physically weak position*” as an instance of *Concept* in the ontology

Figure 4.5 illustrates how the concept “*suspected or accused persons who are in a physically weak position*” is visualised as an instance within the ontology.

Moreover, Figure 4.6 presents the *ImportantExampleDefinition* for this particular concept. Additionally, Figure 4.7 shows the relationship be-

tween this concept and “*suspected or accused persons who have any physical impairments which affect their ability to communicate effectively*”.

IRI	
http://webprotege.stanford.edu/suspected_or_accused_persons_who_are_in_a_potentially_weak_position_EU-exempledefinition	
Annotations	
⚙️ rdfs:label	⚙️ suspected or accused persons who are in a potentially weak position_EU-exempledefinition
⚙️ rdfs:comment	⚙️ Recital 27 in full: The duty of care towards suspected or accused persons who are in a potentially weak position, administration of justice. The prosecution, law enforcement and judicial authorities should therefore ensure that s potential vulnerability that affects their ability to follow the proceedings and to make themselves understood, and
⚙️ source	⚙️ EU Directive 2010-64, Recital 27
⚙️ value	⚙️ suspected or accused persons who are in a potentially weak position, in particular because of any physical impa
Enter property	Enter value
Types	
🔵 ExampleDefinition	
Enter a class name	

Figure 4.6: ImportantExampleDefinition for the concept of “*suspected or accused persons who are in a physically weak position*”




Figure 4.7: Illustrating the connection between “*suspected or accused persons who are in a physically weak position*” and “*suspected or accused persons who have any physical impairments which affect their ability to communicate effectively*”

In summary, Figure 4.8 displays the comprehensive set of definitions in the analogical lightweight ontology. Here, classical definitions are highlighted in red and analogical definitions in green. Definition classes form the framework of the ontology, while the specific instances (recitals and articles) encapsulate where these definitions or their parts are incorporated. Red and green edges were used to indicate their presence in the dataset. Notably, analogical definitions predominate in these six directives, underscoring their importance for legal conceptual modelling aimed at future semi-automatic extraction methods in certain domains.

As discussed, the directives analysed do not contain specific clauses defining the most critical concepts. Therefore, classical definitions are




Figure 4.8: Graphical overview of definitions within the analogical lightweight ontology

compiled from fragments found across various directives, judgments, and other referenced legal texts. Occasionally, as shown in the diagrams, it is not feasible to assemble full classical definitions due to a lack of comprehensive information, such as a lack of purpose-related details. In such cases, a separation was maintained for analogical definitions which, while less complete than classical definitions, still provide crucial information like examples and contextual details related to time or individuals.

4.2.4 Conclusions and future work

The research described above was developed for the *CrossJustice* project, which dealt with six directives characterised by a scarcity of classical definitions and a predominance of various non-classical types, such as Example,

Non-Example, Purpose, and Parameter Definitions. These are instrumental for analogical, teleological, and systematic reasoning. Consequently, a lightweight ontology was tailored to facilitate the computational management and machine-readability of this diverse information, presenting a pioneering approach of benefit to both the *CrossJustice* project and the broader field of legal informatics.

In the methodology, principles are conceptualised and often connected through a *Purpose* link. Notably, principles like “*the right to access to a lawyer*” are typically grounded in higher authority legislation, which the directives reference explicitly. Looking ahead, more detailed modelling of the connections between legislative sources could be explored. For example, several directives implement requirements from international treaties or conventions, such as the ECHR and the CFREU, and often reference policy initiatives like the Stockholm Programme established in 2009. Future analyses might evaluate how effectively these directives achieve the goals set out in EU policies, potentially employing other ontological strategies and network analysis.

In developing this ontology, numerous references to significant legislative texts were noted, such as the Charter of Fundamental Rights of the European Union (CFREU), the European Convention on Human Rights (ECHR), and the International Covenant on Civil and Political Rights (ICCPR). This highlights the complexity of harmonisation within this legal domain, which extends beyond the conventional national and EU jurisdictional framework of the *CrossJustice* project. Future research in legal informatics could address the broader multilevel, international and European protection of fundamental rights. The ICCPR, effective from 1976 and ratified by all EU member states, obliges the EU to adhere to its principles. The dynamic is further complicated by the interplay between the CFREU and the ECHR, particularly under Article 52(3) of the CFREU, which mandates that the interpretation of rights by the ECHR sets a minimum standard that should not be weakened by EU law. The references of the CFREU into EU legislative acts, which are then transposed into

member state laws, and their interpretation by the CJEU underscore the intricate relationship between these texts. Adjudication on the ICCPR is managed by the Human Rights Committee, a panel of independent experts monitoring the covenant’s implementation, the ECHR is overseen by the European Court of Human Rights within the Council of Europe framework, and the CFREU by the Court of Justice of the European Union.

Many scholars are examining the interplay between these legal instruments and their respective judicial or quasi-judicial bodies, a dialogue enriched by a diverse array of legal sources including judgments, decisions, legislation, and non-binding legal and political instruments (soft law). This complex interrelationship, emblematic of multilevel human rights protection, poses significant challenges and opportunities for legal informatics and harmonisation. The domain-specific legal analogical ontology should comprehensively incorporate definitions from all these jurisdictional levels to reflect the multifaceted nature of legal interpretation and application.

4.3 A harmonised glossary for LLM-based legal concept detection

Analysis of transposition deficiencies within the EU, as exemplified by the InfringEye report [9], reveals a substantial number of infringement proceedings initiated by the European Commission due to ineffective transposition of EU legislation into national legal systems. By April 2021, the Commission had opened 1,821 such procedures, with the highest incidence in the environment, energy, mobility, internal market, and justice sectors [334].

Infringement proceedings stem primarily from treaty and regulation violations, delayed directive transposition, and inaccurate directive transposition or application [335]. This chapter presents a methodological framework to assist EU member states in rectifying and preventing such errors.

The complexity of this research reflects the intricate nature of EU legal harmonisation [14]. The persistent challenges highlight the need for innovative legal technology solutions to improve the efficiency and accuracy of conformity assessments. A crucial obstacle lies in ensuring consistent interpretation and application of legal concepts across diverse jurisdictions despite linguistic and cultural variations. Effective transposition hinges on numerous factors, including linguistic complexity, the multilingual nature of EU law, the type of transposing legislation, and the impact of national legal traditions [169]. The integration of semantic knowledge modelling with advanced transformer-based large language models (LLMs) offers a promising avenue for addressing these challenges [336].

The second part of the chapter proposes a Harmonised Glossary (HG), a novel interconnected taxonomic resource encompassing European and corresponding national legal concepts to facilitate analysis of legislative harmonisation as a textual phenomenon and through the lens of legislative approximation. The HG is intended to standardise terminology, thereby promoting consistency and clarity across legal systems. Its focus on fine-grained legislative concepts is intended to address subtle inter- and intra-linguistic discrepancies that impede effective harmonisation.

Utilising a manually-constructed glossary of concepts derived from Council Framework Decision 2002/584/JHA (the European Arrest Warrant Framework Decision), a language model was employed for the semi-automated extraction of corresponding domestic legal concepts, leveraging its inherent linguistic and semantic capabilities. This methodology, incorporating detailed legal annotations, introduces semi-automation into cross-domain legislative text analysis while retaining expert review and validation. This combined semantic and syntactic approach enables comprehensive, nuanced, and interpretable analysis of legislative approximation within the EU and its member states.

HG, coupled with LLM-based concept extraction, provides policymakers and legal professionals with a scalable method of assessing the efficacy of EU legislation and identifying areas requiring further harmonisation.

This research can contribute to the ongoing transposition of EU law into national legislation with the aim of reducing the incidence and severity of infringement proceedings. The ultimate goal is to foster a more coherent, unified, and efficient legal framework across the European Union.

This work details the creation of a graph-based resource that taxonomically links European legislative concepts, augmented with corresponding legislative sources from both the EU and member states. A legal technology expert specialising in legislative harmonisation analysis manually extracted and organised the European concepts. The study also utilises LLMs for the semi-automated identification of European legislative concepts and the corresponding textual portions where they occur within the implementing legislation of five EU member states.

The data underlying this research, including HG, the manually extracted European concepts from Council Framework Decision 2002/584/JHA, and the output of the automated national concept extraction with manual annotations are publicly available at [this GitHub repository](#).

4.3.1 Dataset

This study employs a manually-created dataset, designed for both manual and computational analysis, developed by EU criminal law experts within the EU-funded CrossJustice and Facilex projects. This dataset comprises tables of EU legislation and their national transpositions, systematically segmented into articles and paragraphs. To ensure consistent analysis, domestic legislation in multiple languages has been translated into English. Specifically, the dataset contains official English versions of EU legislative acts alongside the English translations of corresponding NIMs. The JSON format of the data enhances its accessibility for machine processing. This aligns with the EU's interoperability objectives as outlined in the EU Interoperable Act (Regulation 2024/903). The resource includes 32 articles and their national implementations (where available), totalling 144 article-implementation pairs in English from the following countries: France (27),

Germany (26), Italy (31), Portugal (30), and Spain (30).

Figure 4.9 presents a sample from the transposition table, showing Article 1 of the European Arrest Warrant (EAW) Framework Decision segmented into cells (left) and the corresponding legislative portion from the Spanish transposition (right). This visualisation clarifies the structural mapping between the EU framework and national implementations.

Even with detailed transposition tables, significant pre-processing is necessary for effective computational analysis. For instance, textual similarity metrics have been used to assess semantic correspondence between EU and national legislation [337]. This crucial stage in legislative harmonisation analysis intentionally excludes additional legal sources, such as case law, focusing on the mere legislative harmonisation.

Previous computational approaches [338] to legal harmonisation analysis used transposition tables to compare EU legislation with national measures but did not address the challenges posed by high semantic similarity within the same legislative domains. Specifically, identifying domestic measures corresponding to the FD EAW is complicated by the presence of other criminal and procedural law legislation within the corpus.

To overcome the limitations of simple vector-based similarity methods, the research described in this section employs a more nuanced approach, extracting concepts and definitions from the transposition tables. This manual refinement enables the semi-automated retrieval of concepts, definitions, and corresponding implementing articles in order to support more sophisticated computational analyses and precise comparative legal studies. The integration of large language models enhances scalability and enables the models to augment the dataset by retrieving domestic concepts and definitions based on an interconnected, taxonomic glossary of European concepts.

Such analyses are crucial for understanding implementation patterns and discrepancies across legal systems and thus provide insights into the effectiveness of EU legislative harmonisation.

Article 1 Definition of the European arrest warrant and obligation to execute it	
1. The European arrest warrant is a judicial decision issued by a Member State with a view to the arrest and surrender by another Member State of a requested person, for the purposes of conducting a criminal prosecution or executing a custodial sentence or detention order. 2. Member States shall execute any European arrest warrant on the basis of the principle of mutual recognition and in accordance with the provisions of this Framework Decision. 3. This Framework Decision shall not have the effect of modifying the obligation to respect fundamental rights and fundamental legal principles as enshrined in Article 6 of the Treaty on European Union.	Explicitly transposed + Add National act Paragraphs 1 and 2 of this precept are reflected in Articles 34 and 1 respectively of the Spanish law on mutual recognition of criminal decisions. Art. 34 European arrest warrant and surrender. A European arrest and surrender warrant is a judicial decision handed down in a Member State of the European Union with a view to arrest and surrender by another Member State of a person who is claimed to take criminal actions against him or to enforce a custodial sentence or measure of deprivation of liberty, or a measure of internment in a centre for minors. Article 1. Mutual recognition of decisions on criminal matters in the European Union. In application of the principle of mutual recognition of decisions on criminal matters in the area of freedom, security and justice of the European Union, the Spanish judicial authorities who hand down an order or a decision comprised within the provisions of this Act, may transmit it to another Member State for its recognition and execution. In application of the principle of mutual recognition of decisions on criminal matters, the competent Spanish judicial authorities shall recognise and execute European orders and decisions on criminal matters foreseen in this Act within the term foreseen, when they have been transmitted correctly by a competent authority of another Member State and no established ground to refuse recognition or execution concurs. + Add Provision Article 3 - Fully implemented ✓ This Act shall be applied respecting the fundamental rights and liberties and the principles set forth in the Spanish Constitution, in Article 6 of the European Union Comment Ley 23/2014, de 20 de noviembre, de reconocimiento mutuo de resoluciones penales en la Unión Europea Law on mutual recognition of criminal decisions in the European Union + Add National act

Figure 4.9: Excerpt of a transposition table for Article 1 of the FD EAW and corresponding Spanish implementing legislation

4.3.2 Methodology

The HG graph facilitates efficient identification of EU legal concepts, their national equivalents, and definitions. It comprises a network of EU concepts, articles and legislative references manually extracted by legal experts from the FD EAW. These concepts are interconnected within a comprehensive, taxonomic graph structure. This initial phase establishes a foundation for the semi-automated integration of corresponding concepts from national transposing legislation, creating a data representation suitable for evaluating the extent of legislative harmonisation across the EU. The following section details the LLM-based methodology used to automate this harmonisation analysis.

Manual concept identification

A legal technology expert specialising in legislative harmonisation analysis meticulously annotated semantically significant terms within the FD EAW, following a thorough review. This process began with the first normative provision, i.e. excluding recitals. While recitals offer important interpretative context, they are not legally binding [339] and are not subject to domestic transposition, hence their exclusion from this analysis. A concept was considered relevant if its absence would materially alter the meaning of the law or a paragraph and it was specific enough to be crucial for understanding the law’s essential elements.

Figure 4.10 illustrates the concept annotation of Article 1 of the FD EAW. The “European Arrest Warrant (EAW)” concept is fundamental as it concerns the legislative act’s overall objective. Similarly, the principle of mutual recognition of judgments and judicial decisions, underpinning European judicial cooperation in criminal matters (Article 82 TFEU), is essential to the EAW mechanism. Conversely, “entry into force” is deemed to be semantically irrelevant. It is a standard legislative element specifying temporal application and it lacks the semantic weight to alter the meaning of provisions, despite its importance for the law’s validity.

CHAPTER 1

GENERAL PRINCIPLES

Article 1

Definition of the European arrest warrant and obligation to execute it

1. The **European arrest warrant** is a judicial decision issued by a Member State with a view to the arrest and surrender by another Member State of a requested person, for the purposes of conducting a criminal prosecution or executing a custodial sentence or detention order.
2. Member States shall execute any European arrest warrant on the basis of the **principle of mutual recognition** and in accordance with the provisions of this Framework Decision.
3. This Framework Decision shall not have the effect of modifying the obligation to respect fundamental rights and fundamental legal principles as enshrined in Article 6 of the Treaty on European Union.

Figure 4.10: Manual identification of concepts under Article 1 of the Framework Decision on the European Arrest Warrant

Concepts such as “terrorism”, “human trafficking” and “corruption” are relevant to the FD EAW. These offences can, under certain conditions and absent other grounds for refusal, lead to the surrender of a requested person by one member state’s authorities to another, even if the offence is not defined identically in the executing member state’s legal system. The omission of any of these offences would significantly alter the meaning of Article 2 of the FD EAW.

Similarly, “fundamental rights” and “fundamental legal principles” are included in the glossary, per Article 1(3) of the FD EAW. This highlights their crucial role, emphasising that EAW operational provisions must not violate obligations concerning fundamental rights under EU law. This

provision establishes the EAW’s subordination to overarching EU treaties, ensuring that all EAW actions comply with EU fundamental rights. This integration, aligning its application with fundamental European legal principles, strengthens the EAW’s legal foundation and safeguards against misuse.

The inclusion of terms not explicitly defined in the FD EAW, along with implicit concepts, is aimed at comprehensively modelling the law’s essential semantic components at a high level of granularity. This approach mitigates the risk of omitting relevant concepts and significantly expands the annotated corpus. Mapping both explicit and implicit concepts enhances the semantic modelling of the text, resulting in a more detailed and accurate interconnected glossary, thereby improving both precision and usability.

The identification of concepts and relationships required approximately thirty-two hours, including eight hours for revising the initial draft.

Interconnecting concepts in a taxonomic structure

The second phase of analysis employs a taxonomic framework to organise and link related concepts. This structure links concepts to broader categories, enabling accessibility for human experts and natural language processing (NLP) applications, particularly for semi-automated categorisation. These overarching categories, foundational to the taxonomy, are designated as “implicit concepts”. As demonstrated in the next section, these implicit concepts are also integrated into the system prompt given to the language model. This methodological step enhances the identification of both explicit and implicit concepts, as shown in Figure 4.11. Within the HG, concept definitions correspond to their textual locations within the legislative act. Importantly, the only explicit definition of the European Arrest Warrant (EAW) was manually added to the dataset. A legal technology expert categorised concepts and annotated their corresponding article and paragraph references to support semi-automated definition

2. The following offences, if they are punishable in the issuing Member State by a custodial sentence or a detention order for a maximum period of at least three years and as they are defined by the law of the issuing Member State, shall, under the terms of this Framework Decision and without verification of the double criminality of the act, give rise to surrender pursuant to a European arrest warrant:

- participation in a criminal organisation,
- terrorism,
- trafficking in human beings,
- sexual exploitation of children and child pornography,
- illicit trafficking in narcotic drugs and psychotropic substances,
- illicit trafficking in weapons, munitions and explosives,
- corruption,
- fraud, including that affecting the financial interests of the European Communities within the meaning of the Convention of 26 July 1995 on the protection of the European Communities' financial interests,

Figure 4.11: Manual identification of concepts under Article 2(1) of the Framework Decision on the European Arrest Warrant

retrieval.

Figure 2.2 shows that the excerpt from Article 1(2) of the Framework Decision on the European Arrest Warrant (FD EAW) lists offences that, without double criminality verification and subject to other conditions, mandate the surrender of the requested person. This list is crucial to the judicial cooperation mechanism, obligating EU member states to surrender individuals charged with listed offences, regardless of double criminality, unless other grounds for refusal apply. These concepts are grouped under the broader category of “offenses mandating surrender under the European Arrest Warrant”. This categorisation enhances the semantic model of the

FD EAW, simplifies consultation, and introduces the broader taxonomic categories of implicit concepts. This approach also assesses the language model’s capacity for semi-automated retrieval of implicit concepts within the legislative text.

In summary, this process yielded two tabular datasets. The Concept Table (Table 4.6) lists identified concepts and their locations within the FD EAW, providing a basis for semi-automated definition and relationship extraction and indexing. The Concept-Relationship Table (Table 4.7), detailing binary relationships between expert-identified concepts, is used, alongside the Concept Table. This data structure allows for both topological and semantic interconnection of concepts.

A knowledge base for LLM-based concept-detection

While a legal technology expert manually extracted concepts from the FD EAW, an LLM-based approach was used for the corresponding domestic implementing legislation. Creating the initial Hierarchical Glossary (HG) was a labour-intensive manual process. Extending this to all five member states in the dataset would have required a comparable time investment. Consequently, an automated approach, leveraging the initial glossary as a knowledge base, was adopted for extracting EU legal concepts from national implementations.

Close collaboration between legal and computer science experts shaped both the system prompt and the prompt design, with the aim of eliciting the necessary legal knowledge and identifying optimal methodologies. Appendix 4.4 details the system prompt, which covers role definition, resource awareness, context setting, example provision, and goal specification. The LLM is framed as a European legislation and member state implementation expert, cognisant of a pre-existing ontology of European legal concepts. The instructions specify operation within the context of the Framework Decision on the European Arrest Warrant, providing a concrete example of European legislative text and its constituent concepts. Finally, the LLM

Article	EU concepts in long form	EU concepts in short form
Art 1(1), 2(1)	European arrest warrant	European arrest warrant
Art 1(3)	Fundamental rights ex Article 6 TUE	Fundamental rights
Art 1(3)	Fundamental legal principles ex Article 6 TUE	Fundamental legal principles
Art 2(1)	Act punishable by the law of the issuing member state	Punishable act
Art 2(1)	Custodial sentence for a maximum period of at least 12 months	12-month max custodial sentence
Art 2(1)	Detention order for a maximum period of at least 12 months	12-month max detention order
Art 2(1)	Sentence of at least four months	Sentence of at least four months
Art 3(3)	Passing or making a detention order	Detention order issuance
Art 1(1)	Arrest and surrender of a requested person	Arrest and surrender
Art 2(2)	Offence giving rise to surrender pursuant to an European Arrest Warrant	Offence for surrender
Art 4(3)	Participation in a criminal organisation	Criminal organisation
Art 2(2)	Terrorism	Terrorism
Art 2(2)	Trafficking in human beings	Human trafficking
Art 2(2)	Sexual exploitation of children and child pornography	Child exploitation
Art 2(2)	Illicit trafficking in narcotic drugs and psychotropic substances	Narcotics trafficking

Table 4.6: Extract from the Concept Table, which comprises manually annotated concepts and links to legislative articles

is directed to utilise transposition tables to analyse European harmonisation and identify correspondences between European and member state legislation.

For this task, several prominent large language models were considered, distinguished by their scale (parameter count and training data size) and research prominence. Preliminary evaluations involved LLaMA-3.1-405B [340], GPT-4-Turbo [312], Claude-3.5-Sonnet-200k [315], and

EU long-form concepts	Relation	EU long-form concepts
European arrest warrant	is issued by	Issuing judicial authority
European arrest warrant	is issued for	Act punishable by the law of the issuing member state
European arrest warrant	has mandatory compliance with	Fundamental rights ex Article 6 TUE
European arrest warrant	has mandatory compliance with	Fundamental legal principles ex Article 6 TUE
Act punishable by the law of the issuing member state	is punished by	Detention order for a maximum period of at least 12 months
Sentence of at least four months	has condition	Passing or making a detention order
European arrest warrant	is executed by	Executing judicial authority
European arrest warrant	has purpose	Arrest and surrender of a requested person
European arrest warrant	has mandatory execution for	Offence giving rise to surrender pursuant to an European arrest warrant
Offence giving rise to surrender pursuant to an European arrest warrant	has instance	Sexual exploitation of children and child pornography
Offence giving rise to surrender pursuant to an European arrest warrant	has instance	Illicit trafficking in narcotic drugs and psychotropic substances

Table 4.7: Extract from the manually annotated Concept-Relation Table, which represents relations between legal concepts. The directed relations in the second column are from the concepts in the first column to the concepts in the second column.

Gemini-1.5-Pro [314]. Claude-3.5-Sonnet-200k was chosen for its superior accuracy and reduced tendency towards factual inaccuracies (“hallucinations”), demonstrating a preference for non-response over misleading outputs.

The prompt engineering strategy employed was the Tabular Chain of Thought method [341, 342]. This approach encourages a stepwise reason-

ing process involving intermediate calculations before reaching the final objective. The tabular output format facilitates the alignment of European and corresponding domestic concepts for comparative analysis.

The extraction prompt is detailed in Appendix 4.5. It outlines European legal harmonisation and the legal concept extraction task. To specify the desired output format, it defines the following table columns: *European concept* (a manually annotated concept from Harmonised Glossary), *corresponding national concept* (a model-identified related national concept), *nature of concept* (“explicit” or “implicit”), *type of concept match* (“semantic exact match”, “semantic partial match”, or “not applicable”), *corresponding national definition* (a list of national implementation excerpts containing the identified concept), and *notes* (free-text explanatory notes). Two examples, illustrating explicit and implicit concepts, are provided. The prompt then instructs the model to generate the table, given a list of European concepts and their potential locations within national legislative implementations. Finally, standard markdown notation specifies the desired tabular output format: “— European concept — corresponding national concept — nature of concept — type of concept match — corresponding national definition — notes —”. This approach leverages the model’s capacity to identify corresponding national concepts using both the text provided and expert-extracted concepts.

A key feature of the prompt is its inclusion of both explicit and implicit concepts. This assesses not only the model’s ability to align explicitly defined concepts (including synonyms) but also its capacity to infer implicit content primarily representing the taxonomic categories underlying the glossary structure. While not explicitly stated in the prompt, the largely hierarchical relationships between explicit and implicit concepts seem to be beneficial for task completion.

As detailed in the next section, the model demonstrates high accuracy in retrieving and aligning both explicit and implicit concepts. The resulting concept table includes columns indicating whether a concept is explicit or implicit, the degree of semantic match (partial or total) between Euro-

pean and national concepts, the location of the concepts in the legislative text , and it provides qualitative analysis of each alignment within the “notes” column.

Evaluating LLM-based concept detection

This subsection details the methodology for assessing the model’s capacity for semi-automated legal concept extraction from legislative texts. Appendix 4.6 provides the detailed annotation guidelines, thereby ensuring experimental reproducibility.

Two evaluation approaches were initially considered: multiple annotators versus a single annotator. A single-annotator approach was selected, entrusting validation to the expert who had compiled the initial FD EAW glossary. This deviates from standard practice but offers a more focused approach given the specificity of the context and task. The annotator’s expertise ensured precise identification of both explicit and implicit concepts (the latter often representing taxonomic categories) and consistent assessment of total or partial semantic correspondence. The expert also efficiently linked concepts to their definitions as extracted from the legislation. While using a single annotator is unconventional, this choice is justified by the unique characteristics of the evaluation, which enables direct comparison between LLM’s annotations and the annotations of the original glossary creator.

Section 4.3.3 provides an example annotation and validation rationale focusing on accuracy, explainability, and consistency. For transparency and further analysis, the annotation results are publicly available in [this GitHub repository]. The annotation guidelines (Appendix 4.6) describe a binary classification task performed by the annotator on model outputs to assess concept alignment.

Additionally, the annotator performed two further tasks: validating returned concepts and validating omitted concepts. Although the prompt (see Appendix 4.5) requests tabular output, the model might omit speci-

fied European concepts. The second task determines whether these omissions reflect genuine absence from the national implementations or just the limitations of the model. Omissions reflecting genuine absence do not introduce errors into HG.

In **Task 1**, annotators were asked the following questions:

1. *Q1: Is the value returned for the “nature of concept” field correct?*
Yes/No
2. *Q2: Is the value returned for the “type of concept match” field correct?* Yes/No
3. *Q3: Is the value returned for the “corresponding national definition” field correct and comprehensive?* Yes/No

Task 2 consisted of the following question:

1. *Q4: Is the omitted concept present in the national segment?* Yes/No

For questions Q1, Q2, and Q3, a higher proportion of “Yes” responses indicates greater automated system accuracy. Conversely, for question Q4, a higher proportion of “No” responses is preferred, indicating fewer omissions of concepts present in the text. Detailed explanations of the annotation questions are provided in Appendix [4.6](#).

The implementing tool: Neo4J

The selection of Neo4j [\[343\]](#) for representing the interconnected concept graph derived from the FD EAW is motivated by several factors, reflecting its prevalent use in handling complex data structures [\[344\]](#). This graph database is well-suited to the intricate data structures required for mapping the complex relationships between legal concepts and legislative articles.

Neo4j’s robust visualisation capabilities facilitate the interpretation of complex networks of nodes and relationships — a crucial feature in the legal domain where clarity is paramount. Furthermore, its support for complex network queries enables users to explore relationships between legislative concepts, thus aiding in legal interpretation and precedent identification.

Thirdly, Neo4j’s scalability ensures that performance is maintained even with substantial growth in concepts and relationships — a critical consideration for legal research projects with potentially large datasets. Finally, its robust transaction management capabilities are essential for preserving data integrity, given that modifications to concepts or relationships can have significant consequences. Additionally, the graph database supports export in standard ontological formats such as RDF (.rdf, .rdfs, .ttl).

4.3.3 Results

As previously noted, the final glossary comprises manually identified concepts, relationships, and attributes curated by a legal expert. This glossary was subsequently augmented with corresponding national legislative concepts from multiple EU member states. The first subsection presents results from the initial glossary version; the second presents results from the complete, meticulously validated glossary.

The EU concept-based glossary

The Neo4j graph, constructed from manually extracted FD EAW concepts, comprises 124 articles, 198 EU concepts, 212 “found_in” relationships, and 225 interconcept relationships. Figure [4.12](#) illustrates the overall structure of this graph-based glossary.

The graph’s complexity, characterised by numerous highly connected nodes (hubs), stems from several factors. The granular concept identifi-




Figure 4.12: Concepts from the Framework Decision on the European Arrest Warrant in Hierarchical Glossary

cation results in numerous subclasses of broader concepts, such as those categorising crimes automatically triggering EAW issuance. This structure also optimises information retrieval. The legal expert emphasised the roles of key institutions (issuing and executing national authorities and member state responsibilities) in the EAW process, resulting in highly connected nodes representing these elements. Furthermore, the varying density of concepts across different articles leads to uneven node representation within the graph. The graph-based glossary comprises entities, including nodes and attributes, and their relationships, detailed as follows:

Entities

- **EUconcept**: represents an extracted concept, interconnected through relations manually defined by legal domain experts. Attributes for each **EUconcept** node includes:
 - **id**: item identifier;
 - **name**: textual identifier;
 - **long_name**: the full name of the concept as manually extracted by domain experts;
 - **short_name**: a shortened name to display on the node for the purpose of readability;
 - **definition**: a segment of the EU directive defining the concept (if present);
 - **eu_source**: a list of EU legislative segments manually annotated that contains the EUconcept;
 - **[nation]_source**: for each nation, a list of legislative portions of text automatically extracted from the transposition tables.
- **Article**: specifies the location in which each EUconcept was found, expressed as an article, paragraph, and sub-paragraph (if present). Attributes include:
 - **id**: item identifier;
 - **name**: the full name of the concept as manually extracted by domain experts.
 - **EUlegislation**: Specifies the document to which the article belongs. Attributes include:
 - **id**: item identifier.
 - **name**: the full name of the concept as manually extracted by domain experts.

Relations

- **interconcept relation(s)**: customised relation(s) between concepts obtained after manual annotation;
- **found_in**: links the concept to the related segment of the EU Directive;
- **belongs_to**: links the segment of the EU Directive to the EU Directive itself;

The initial, manually created glossary contained only concepts from the FD EAW; therefore, the [nation]_source field is initially empty. The following subsection provides an example of entities representing national concepts and sources, focusing on the central concept of the European Arrest Warrant (EAW), as graphically illustrated in Figure 4.13.

Figure 4.13 depicts the EAW concept (note that node labels may be truncated for clarity). The EAW concept is linked to Article 1(1) via a “found_in” relationship, reflecting its initial occurrence. Article 1(1) is further linked to the entire FD EAW via a “belongs_to” relationship. Finally, the EAW concept is linked to “arrest and surrender of the requested person” via a “has_purpose” relationship, highlighting the EAW’s primary objective: the executing state’s arrest and surrender of the requested individual to the issuing state.

Unlike the relatively straightforward representation of the EAW concept (Figure 4.13), the depiction of the criminal offences listed in Article 2(2) of the FD EAW (Figure 4.14) is more complex. Each criminal offence (represented as a concept) is linked to the “offence giving rise to surrender” node, a taxonomic category encompassing all the offences listed in Article 2(2). These offences are linked via “has_instance” relationships, indicating their instantiation of this broader category. Each offence concept is also linked to the article containing that concept (orange node), which in turn is linked to FD EAW (green node).




Figure 4.13: The modelling of the concept of European arrest warrant and its relations within the FD EAW

LLM-based retrieval of national concepts

The graph structure remains consistent despite the addition of national concepts corresponding to the EU concepts. However, the number of node attributes has significantly increased, as illustrated below for the EAW concept. The EAW concept, refers to the one explicit definition in the legislation. For each EU member state, corresponding national legislative measures containing the concept are also included. This approach ensures that the graph accurately reflects variations in EAW implementation across different national legal systems. The EAW concept attributes are presented below:

- **Name:** European Arrest Warrant
- **Long Name:** European arrest warrant
- **Short Name:** EAW
- **Definition:** The European arrest warrant is a judicial decision issued by a member state with the aim of the arrest and surrender by another member state of a requested person, for the purposes of

- **Source from Italy:** Execution of the European arrest warrant must not violate specific conditions. It shall be executed regardless of double criminality for offences punishable by at least three years as described in Article 2(2) of the Framework Decision.
- **Source from Portugal:** The European arrest warrant may be issued for acts punishable by a custodial sentence or a detention order for at least twelve months, or for sentences of at least four months.
- **Source from Spain:** When a European arrest and surrender warrant has been issued.

Beyond the development of HG, analysing the model’s output across various prompt iterations revealed Claude-3.5-Sonnet-200k’s ability to: identify corresponding national concepts from EU concepts, distinguish between explicit and implicit concepts, detect total or partial semantic correspondence, locate the concept within the text, and provide explanatory notes on the EU-national concept overlap. This detailed output highlights the model’s sophisticated handling of complex legal terminology across diverse legislative frameworks.

Table 4.8 exemplifies Claude-3.5-Sonnet-200k’s output, focusing on the alignment of four EU concepts with their French counterparts. The EAW concept, for example, shows an exact, explicit match. Conversely, “detention order for a maximum period of at least 12 months” lacks a direct French equivalent, demonstrating an implicit concept potentially inferable from other provisions. Concepts like “illicit trafficking in narcotic drugs and psychotropic substances” and “computer-related crime” exhibit partial matches, with French legislation omitting qualifiers or using equivalent terminology (e.g. “cybercrime”). “Environmental crime, including the trafficking of endangered species” shows a substantive, albeit terminologically nuanced, equivalence. The model also identifies corresponding national definitions and their locations in the legislative texts. The “notes” column notably showcases the model’s reasoning capabilities in analysing

the degree of semantic overlap. The examples in Table 4.8 represent the overall semi-automated concept retrieval and alignment results.

EU concept	Corresponding national concept	Nature of the concept	Type of concept match	Corresponding national definition	Notes
european arrest warrant	European arrest warrant	explicit	semantic exact match	["Under the law of the issuing member state, the following offences may give rise to a European arrest warrant:", "The execution of a European arrest warrant may also be refused if the offence for which the warrant is issued does not constitute an offence under French law."]	The concept is directly mentioned in the national legislation.
detention order for a maximum period of at least 12 months	none	implicit	not applicable	[]	The concept is not explicitly mentioned in the national legislation, but may be implied in the custodial sentence provisions.
illicit trafficking in narcotic drugs and psychotropic substances	Trafficking in narcotic drugs and psychotropic substances	explicit	Semantic partial match	["5° Trafficking in narcotic drugs and psychotropic substances;"]	The national legislation omits the word "illicit, but the meaning is implied.
computer-related crime	Cybercrime	explicit	Semantic partial match	["11° Cyber-crime;"]	The national legislation uses a more modern term, but the concept is essentially the same.
environmental crime, including illicit trafficking in endangered animal species and in endangered plant species and varieties	Crimes and offences against the environment, including trafficking in endangered animal species and endangered plant species and essences	explicit	Semantic partial match	["12° Crimes and offences against the environment, including trafficking in endangered animal species and endangered plant species and essences;"]	The national legislation uses slightly different wording but covers the same concept.

Table 4.8: Alignment of EU and national corresponding concepts from Article 2 of the Framework Decision on the European Arrest Warrant and the French implementing legislation

Figure 4.13 presents the results. The model accurately identified identified explicit and implicit concepts in 98% of cases. This high accuracy is partly attributable to the use of the same annotator for both the initial manual HG creation and the subsequent assessment, which minimised subjective interpretation bias. The model effectively distinguished between explicitly defined concepts and those serving as taxonomic categories within




Figure 4.15: Results of the annotator’s responses to questions Q1, Q2, and Q3 on validating returned concepts and validating omitted concepts.

the glossary’s semantic structure.

However, performance in identifying the type of concept match was less consistent and appeared task-dependent. The instruction to distinguish between full and partial semantic matches may have been insufficiently precise. While the model identified semantic overlaps and discrepancies, providing descriptions in the notes, the results lacked consistent legal rigour. Its performance in simply indicating the degree of semantic overlap proved to be inconsistent and lacking in sufficient explanatory detail.

To address this limitation, the annotator refined the partial match categorisation. Partial matches potentially leading to inconsistent legal interpretations across jurisdictions were classified as positive instances; mere semantic differences were classified as negative instances (and thus considered as full matches). For example, “expenses of the executing member state” (Article 30 FD EAW) showed a partial match with the Portuguese “expenses incurred in the national territory”. The Portuguese phrasing, while adapted to its context, does not lead to legally divergent interpretations.

Conversely, “possible revocation of the right to consent to surrender” (Article 13 FD EAW) showed a partial match with the Portuguese concept

of “consent to surrender may not be revoked”. The Portuguese phrasing explicitly negates revocation, differing significantly from the EU version. In summary, the model consistently identified the correct textual segments for each concept, confirming the accuracy of the initial manually-created HG.

4.3.4 Conclusions

This study demonstrates Claude 3.5 Sonnet 200k’s significant capacity for extracting and aligning legislative concepts from national implementations with their EU counterparts. These findings represent substantial progress towards semi-automatically aligning European legislative articles with national implementations — a crucial aspect of computationally analysing legal harmonisation within the European Union.

Future research, maintaining a multidisciplinary approach, will focus on enhancing multilingual legal analysis. A key step involves applying this methodology to a multilingual dataset and rigorously evaluating the results. This includes expanding HG to facilitate broader cross-lingual comparisons, moving beyond traditional translation-based methods. The expanded HG will enable the exploration of concept-based similarity measures, identifying conceptual convergence or divergence across jurisdictions regardless of whether or not direct translations are available. Further exploration of LLM-based semi-automated concept extraction, benchmarked against a manually curated “gold standard” across diverse legislative areas, will refine the HG and ensure its accuracy. These advancements will provide researchers and legal professionals with a more nuanced understanding of legal harmonisation in multilingual environments such as the European Union.

In conclusion, analysing legal harmonisation requires considerations beyond semantic-textual matching of legislative texts; the effectiveness of implementation also depends on non-legislative sources, such as case law. While methodologies exist for network analysis and semantic modeling of

case law, integrating these with LLMs represents a significant, yet unexplored, research area.

While Hicks et al. [345] contend that LLMs produce unreliable outputs due to their inherent disregard for factual accuracy, this perspective neglects the potential for effective LLM application through appropriate prompting, grounding techniques, and human oversight. Their focus on LLMs' tendency to generate plausible but inaccurate information overlooks emerging research demonstrating that careful guidance and integration with external knowledge sources can significantly improve LLM output reliability and utility. Therefore, the assertion that LLMs are inherently unreliable is an oversimplification that ignores the evolving capabilities of LLMs and their potential for productive application.

Conclusions

The research presented in this thesis deepens the study of the phenomenon of legal harmonisation within the European Union, highlighting its inherent complexity, which is influenced by linguistic, legal, cultural and political factors. Despite the challenges, with the adoption of cutting-edge language models, specifically LLMs, semi-automated legal harmonisation analysis can be achieved. While the methodologies presented focused primarily on the legislative textual analysis, comprehensive understanding requires acknowledging the broader context of the harmonisation process itself. Accordingly, textual analysis is not presented in isolation but rather situated within the rich complexity of the legal landscape. The integration of diverse legal tech methodologies, particularly the use of transformer-based architectures and cutting-edge language models in combination with semantic knowledge modelling, facilitated multifaceted analysis of the phenomenon under investigation. This approach yielded a comprehensive and nuanced understanding, overcoming some of the limitations inherent in employing single methodologies.

The first chapter provides a comprehensive analysis of EU legal harmonisation, moving beyond traditional doctrinal approaches to highlight the intricate interplay between several factors. The inherent complexities of multilingualism within the EU framework are thoroughly explored, referencing relevant treaties and regulations like EU Regulation No. 1/1958 to demonstrate the potential for linguistic discrepancies. The chapter meticulously details the EU legislative procedure, including the ordinary legislative procedure and codecision, and analyses specific cases to illustrate the

concrete challenges of adopting and implementing EU legislative acts. Furthermore, a detailed examination of the CJEU’s interpretive methodologies — literal, systematic, teleological, and comparative — underscores the dynamic and evolving nature of EU law and its impact on harmonisation efforts. By combining these legal and linguistic perspectives, the chapter establishes the need for innovative, technologically-driven approaches to address the multifaceted challenges of achieving legal harmonisation and, possibly, uniformity across member states.

The second chapter addresses the application of legal technology, specifically similarity metrics, to analyse legal harmonisation, with particular attention literal and comparative interpretation. While this thesis has a technological focus, the second chapter effectively demonstrates both the potential and limitations of using traditional NLP techniques for assessing semantic proximity in legislation. The application of vector-based approaches, such as TF-IDF for vector extraction and cosine similarity for comparison, within the context of the EU-funded CrossJustice project, proved to be effective for aligning European directives with corresponding national legislative measures as they offer a quantifiable way to evaluate harmonisation. However, the chapter also implicitly highlights the limitations of relying solely on these traditional methods. To address these limitations, two novel similarity measurements were developed within the Facilex project — an asymmetrical similarity metric and a concept-based similarity metric. The latter, leveraging the harmonised glossary detailed in the fourth chapter, significantly enhances the precision of legislative alignments by grounding comparisons in shared conceptual understanding rather than just lexical overlap. This refinement, demonstrated through case studies, could simplify the transposition of EU directives into national legal systems. The chapter shows that legal technology has the potential to improve the efficiency and accuracy of harmonisation efforts while simultaneously acknowledging the need for more sophisticated approaches beyond simple lexical comparison. The chapter’s findings thus serve as a crucial first analytical step, demonstrating the value of enhanced, concept-driven

methodologies for achieving a more nuanced and accurate assessment of legislative alignment.

The third chapter identifies network analysis as a novel approach to understanding EU legal harmonisation, moving beyond traditional textual analysis to reveal implicit connections and contextual nuances. The chapter illustrates this with three distinct case studies, each highlighting the methodology's unique capabilities. One case study focuses on EU health law, employing graph-based NLP to construct complex networks representing relationships between legal norms, thereby uncovering hidden patterns. A second case study addresses the interconnection between recitals and articles within EU legislation, using co-occurrence networks to highlight implicit connections that significantly impact contextual interpretation. This involved a rigorous annotation process and use of the card sorting methodology to group similar norms with the aim of ultimately revealing the underlying structure and logic of the legislation. Finally, a CJEU case study utilises network analysis to map judicial citation patterns, providing insights into the evolution of legal interpretation and the influence of key precedents across member states. This last study particularly highlights the capacity of network analysis to integrate jurisprudential sources, which are crucial for understanding the implementation of EU law within national legal systems. In summary, the third chapter argues that network analysis offers a powerful lens for comprehending the multifaceted nature of EU legal harmonisation, showing hidden relationships and enhancing contextual understanding.

The fourth chapter explores the application of ontological modelling and the creation of a harmonised glossary to address the challenges of multilingualism and analogical interpretation in harmonisation. The chapter presents two key case studies that demonstrate practical application of these methodologies. Firstly, the development of an analogical lightweight ontology for EU criminal procedural rights within the context of judicial cooperation offers a novel approach to identifying and analysing discrepancies between different linguistic versions of European legislative acts.

This ontology facilitates detailed examination of the slight divergences and potential misalignment that can arise during the translation process, providing a clear view of the interpretative challenges inherent in multilingual legal frameworks. The results revealed significant discrepancies between linguistic versions of legislative acts, highlighting how these variations can lead to divergent interpretations of the same law across member states. This underscores the crucial role of ontological modelling in identifying and mitigating the risks of inconsistent application of EU law.

Secondly, the creation of a harmonised glossary for LLM-based legal concept detection represents a significant step towards enhancing the accuracy and efficiency of automated legal analysis. This glossary, manually constructed by legal harmonisation experts, provides a structured and semantically labelled knowledge base that facilitates automated alignment of concepts by advanced LLMs such as Claude-3-sonnet. The chapter details the methodology, from manual concept identification and taxonomic structuring to implementation using Neo4J, a graph database. Preliminary results demonstrate the promise of this approach to ensuring the scalability and reproducibility of methodologies adopted for legal harmonisation, although more comprehensive evaluation is ongoing. The success of this approach hinges on careful crafting of the glossary, reflecting the need for continuous refinement and expansion of the knowledge base to improve its accuracy and reduce interpretative discrepancies among member states. The chapter highlights the potential of this approach for future research, particularly in leveraging the ongoing evolution of DL technologies and the increasing precision of LLMs to further revolutionise the field of European legal harmonisation. The two case studies presented demonstrate the practical application of these methodologies, highlighting their ability to identify interpretative discrepancies, facilitate automated legal analysis, and pave the way for future advancements in the field of legal harmonisation through the effective integration of cutting-edge AI technologies.

To effectively apply the multi-method approach proposed in this thesis, legal professionals should develop a foundational understanding of the

functioning of generative artificial intelligence, with particular emphasis on LLM-based systems. The proposed methodology is grounded in the extraction of legal concepts from European and national legislative acts and integrates a manually curated legal semantics framework developed by domain experts. However, since LLMs do not produce deterministic outcomes and may generate inaccuracies or hallucinations, it is essential that end users remain aware of these limitations and adopt appropriate mitigation strategies.

In this regard, it is crucial that legal practitioners receive comprehensive training on the methodologies underlying the digital tools they employ, especially considering the growing impact of technology on the legal profession. Such training should cover the fundamental principles of LLMs, their capabilities and limitations, as well as techniques for manual validation and result verification to ensure that AI-assisted analyses meet the rigorous standards required in the legal domain. While the integration of grounding strategies and manually crafted legal knowledge enhances the reliability of LLM-generated outputs, human oversight remains indispensable to validate and contextualize findings accurately.

Appendix

4.4 System prompt

You are an expert in European legislation and the corresponding national implementing legislation of member states. You have a deep knowledge of legal concepts that appear both in European and national legislation. Notably, you are familiar with a non-formal ontology of European legal concepts that has been manually developed by legal experts. These experts have linked these concepts to specific segments of European legislation, including the Framework Decision on the European Arrest Warrant (EU Decision 2002/584).

This is an example: European legislation segment: “1. A European arrest warrant may be issued for acts punishable by the law of the issuing member state by a custodial sentence or a detention order for a maximum period of at least 12 months or, where a sentence has been passed or a detention order has been made, for sentences of at least four months.” European legal concepts found: [‘european arrest warrant’, ‘act punishable by the law of the issuing member state’, ‘custodial sentence for a maximum period of at least 12 months’, ‘detention order for a maximum period of at least 12 months’, ‘sentence of at least four months’]

In a European research project, tables have been manually crafted to show the alignment between the provisions of the Framework Decision on the European Arrest Warrant and the corresponding national laws implemented in various European countries. The goal is to automate the

analysis of legislative harmonization across Europe, defined as the lexical and semantic alignment between European legislation and its national implementations. Your task is to identify correspondences between concepts in European laws and their equivalents in national legislation. Given articles and concepts from European laws, you will match them with the appropriate national legislation, noting the presence of identical concepts or their synonyms.

4.5 Prompt for legal concept detection

Create a mapping table of all European concepts manually extracted by legal experts to assess how accurately they have been transposed into national law. More precisely, identify possible concepts in national legislation that correspond to European concepts. Output in this following format: “| european concept | corresponding national concept | nature of concept | type of concept match | corresponding national definition | notes |”.

The “european concept” column contains concepts that are extracted from the European legislative texts. Use the example provided above to identify concepts under scrutiny.

The “corresponding national concept” column can have as a specification either no or the corresponding concept found in the national transposing legislative segments.

The “nature of concept” column can have values either “explicit” or “implicit”. It refers to the presence, if any, explicit or implicit of the concept in national legislation.

The “type of concept match” column could have the following values that explains the type of match between “european concept” and “correspond-

ing national concept” columns:

- *Not applicable*
- *Semantic exact match*: all meanings of the european legal concept are present in the national legislation, differences such as verb conjugations, stopwords, numbers, or other textual components with weak meaning are to be ignored
- *Semantic partial match*: some of the european legal concept meanings are present in the national legislation but substantial differences were omitted from the perspective of legal language, e.g. the presence of “illicit”, “ may”, “ obligation”, “right”, etc.

The “corresponding national definition” column contains a list of national sentences, which are national legislative portions that contain the “corresponding national concept”. A concept could be found in more legislative portions.

The “notes” column contains an explanation of the concept match found.

I give you two examples of table rows.

Example 1:

- “*european concept*”: custodial sentence for a maximum period of at least 12 months
- “*corresponding national concept*”: custodial sentence of one year or more
- “*nature of concept*”: explicit
- “*type of concept match*”: Semantic partial match

- “*corresponding national definition*”: [“2° Offences punishable by a custodial sentence of one year or more or, where a custodial sentence has been pronounced, where the sentence pronounced is four months’ imprisonment or more.”]

Example 2:

- “*european concept*”: offence giving rise to surrender pursuant to an european arrest warrant
- “*corresponding national concept*”: offence giving rise to execution of european arrest warrant
- “*nature of concept*”: implicit
- “*type of concept match*”: Semantic partial match
- “*corresponding national definition*”: [“By way of derogation from the first paragraph, a European arrest warrant shall be executed without verification of the double criminality of the offences for which it is issued when those offences are, under the law of the issuing member state, punishable by a custodial sentence of at least three years’ imprisonment or by a detention order of a similar duration and fall within one of the categories of offences provided for in Article 694-32.”]

Now compute the table given the following European legal concepts and the national legislative segments.

Concepts manually extracted by legal experts in the European legislative segment:

““““

[concept 1, concept 2, concept 3, ...]

””””

Corresponding legislative segments of national implementation in list format:

““““

[national implementation portion 1, national implementation portion 2, national implementation portion 3, ...]

””””

“| european concept | corresponding national concept | nature of concept | type of concept match | corresponding national definition | notes |”.

4.6 Annotation guidelines

This document provides the guidelines for a manual annotation task designed to evaluate the performance of an automated system for detecting European legal concepts within member state legislation. The results of this annotation will be used to assess the potential of this automation to assist policymakers and legal practitioners in evaluating the effectiveness of European legislation.

4.6.1 Task Context

European legislation is implemented at the national level by each member state. To facilitate analysis, a matching table has been created by legal experts. This table links specific articles in European legislation to their corresponding implementing articles within the national laws of each member state. This pre-existing mapping table constitutes the ground truth for national legislative implementations and will be provided to support the annotation and evaluate the automatically generated results.

4.6.2 Task Objective

An automated system has been developed to identify specific legal concepts within legislative texts. Using the expert-created mapping table, your task is to evaluate the output of this automated system. This evaluation will help us understand the strengths and weaknesses of the automated legal concept detection process and its potential utility in real-world scenarios.

4.6.3 Annotation procedure

In this section, you can find information about the inputs, tasks, and tools to use to complete the annotation.

Input

You will be provided with the following information:

1. A specific article from European legislation.
2. The corresponding national implementing articles (as identified by the expert matching table).
3. The automatically generated output table contains the following:
 - (a) *European Concept*: A legal concept manually annotated from the Council Framework Decision 2002/584/JHA on the European Arrest Warrant.
 - (b) *Corresponding National Concept*: The corresponding national concept, if found.
 - (c) *Nature of Concept*: Whether the concept is implicit or explicit.
 - (d) *Type of Concept Match*: Indicates whether the match is a semantic exact match, a semantic partial match, or not applicable.

- (e) *Corresponding National Definition*: Text segments within the national implementing articles where European legal concepts should be present.
- (f) *Notes*: A generated comment explaining the reasoning behind the row's generation.

Task 1

Your task is to assess the accuracy of the automated system's output for each provided pair of European and national implementing articles. Specifically, you will answer three questions for each table row with "Yes" or "No" values:

1. *Q1: Is the "nature of concept" field correct?*
 - *Yes*: after reading the national implementation, determine whether the European concept is present explicitly, implicitly, or not at all. Then, compare your assessment with the automated system's output. If your assessment matches the system's output, the answer is considered correct.
 - *No*: otherwise
2. *Q2: Is the "type of concept match" field correct?*
 - *Yes*: the automatic process identifies European concepts in national implementations as either a "semantic match" (partial or exact) or "not applicable". A "semantic match" is correct if the concept is present; "not applicable" is correct if it's absent.
 - *No*: otherwise
3. *Q3: Is the "corresponding national definition" field correct and comprehensive?*
 - *Yes*: the system's output is considered correct if it accurately identifies the presence or absence of the European legal concept

derived from the matched European article. Furthermore, if the concept is present, the list of corresponding national legislative portions must be accurate and complete for the output to be considered correct.

- *No*: otherwise

Task 2

As the automatic system sometimes ignored some European legal concepts, your task is to answer the following question for those omitted:

1. *Q4: Is the omitted concept present in the national segment?*
 - *Yes*: after reading the national implementation, you found that the concept omitted by the system is actually present in the national implementation
 - *No*: otherwise

Annotation Tool

You will be provided with two spreadsheets, the first for Task 1 and the second for Task 2. Respond with “yes” or “no” in the cells corresponding to the European concept and the question you are answering.

Important Considerations

1. *Ambiguity*: If you encounter any ambiguity in the automated system’s output or the expert mapping, please document it clearly in the space provided for comments.
2. *Consistency*: Maintain consistency in your annotations throughout the task. If you are unsure about a specific case, refer back to these guidelines or consult with the research team.

By carefully following these guidelines, you will contribute valuable data to this research project, ultimately helping us to better understand how automated systems can support the analysis of European legislation and its implementation across member states.

We thank you for participating in this experiment.

Bibliography

- [1] Mark Van Hoecke. “Methodologies of legal research”. In: (2011).
- [2] Alfred F Conard. “The Quantitative Analysis of Justice”. In: *J. Legal Educ.* 20 (1967), p. 1.
- [3] Rossana Deplano. “The use of international law by the United Nations Security Council: An empirical framework for analysis”. In: *Emory International Law Review* 29 (2015).
- [4] Wolfgang Alschner. “The computational analysis of international law”. In: *Research Methods in International Law*. Edward Elgar Publishing, 2021, pp. 203–227.
- [5] Elena Kantorowicz-Reznichenko. “Computational methods for legal analysis: The way forward?” In: *Erasmus L. Rev.* 14 (2021), p. 1.
- [6] Stephen Weatherill. “Harmonisation: How Much, How Little”. In: *Eur. Bus. L. Rev.* 16 (2005), p. 533.
- [7] Irene Otero Fernández. “Multilingualism and the Meaning of EU Law”. PhD thesis. European University Institute, 2020.
- [8] Emilie Ghio. “Harmonisation and European Integration in Times of Crisis”. In: *Are Legal Systems Converging or Diverging? Lessons from Contemporary Crises*. Springer, 2024, pp. 255–276.
- [9] EdjNet and Openpolis. *Understanding EU Infringement Procedures Through Data*. Available online: <https://www.openpolis.it/infringeye-il-monitoraggio-delle-infrazioni-europee-con-i-dati/>. Accessed on [19.09.2024]. 2023.

- [10] Ernesto Rossi and Altiero Spinelli. *Il manifesto di Ventotene*. Edizioni Mondadori, 2017.
- [11] Cornelis JW Baaij. *The role of legal translation in legal harmonization*. Kluwer Law International BV, 2012.
- [12] Marek Stych. “The definition of family in international and EU law”. In: *MEST Journal* 15 (2021).
- [13] Krystyna Slany. *Alternatywne formy życia małżeńsko-rodzinnego w ponowoczesnym świecie*. ZW nomos, 2002.
- [14] Maja Bratanić and Maja Lončar. “The myth of EU terminology harmonization on national and EU level”. In: *Language and culture in EU law*. Routledge, 2016, pp. 207–218.
- [15] Stephen Weatherill. “Maximum versus minimum harmonization: Choosing between unity and diversity in the search for the soul of the internal market”. In: *From Single Market to Economic Union: Essays in Memory of John A Usher* (2012), pp. 175–199.
- [16] Francesco De Cecco. “Minimum harmonization and the limits of Union fundamental rights review: TSN and AKT”. In: *Common Market Law Review* 58.1 (2021).
- [17] Kai P Purnhagen and Justus HH Wessler. “Maximum vs minimum harmonization: what to expect from the institutional and legal battles in the EU on gene editing technologies”. In: *Pest Management Science* 75.9 (2019), pp. 2310–2315.
- [18] Paul Voigt et al. “Practical implementation of the requirements under the GDPR”. In: *The EU General Data Protection Regulation (GDPR) A Practical Guide* (2017), pp. 245–249.
- [19] Doris Liebwald. “Law’s capacity for vagueness”. In: *International Journal for the Semiotics of Law-Revue internationale de Sémiotique juridique* 26.2 (2013), pp. 391–423.
- [20] *Judgment of the Court of 17 April 1986 - State of the Netherlands v Ann Florence Reed*. Apr. 17, 1986.

- [21] *Judgment of the Court of 31 May 2001 - D and Kingdom of Sweden v Council of the European Union*. May 31, 2001.
- [22] Koen Lenaerts. “The European court of justice and the comparative law method”. In: *European Review of Private Law* 25.2 (2017).
- [23] Koen Lenaerts and José A Gutiérrez-Fons. “To say what the law of the EU is: methods of interpretation and the European Court of Justice”. In: *Colum. J. Eur. L.* 20 (2013), p. 3.
- [24] *Judgment of the Court of 12 November 1969 - Erich Stauder v City of Ulm - Sozialamt*. Nov. 12, 1969.
- [25] *Judgment of the Court of 6 October 1982 - Srl CILFIT and Lanificio di Gavardo SpA v Ministry of Health*. Oct. 6, 1982.
- [26] Alexander Kornezov. “The new format of the acte clair doctrine and its consequences”. In: *Common Market L. Rev.* 53 (2016), p. 1317.
- [27] Bruno De Witte. “Language Law of the European Union: Protecting or Eroding Linguistic Diversity?” In: (2004).
- [28] Jacques Ziller. “Multilingualism and its Consequences in European Union Law”. In: *Common European Legal Thinking: Essays in Honour of Albrecht Weber* (2015), pp. 437–456.
- [29] Diana Yankova. “The text and context of EU directives: implications for translators”. In: *Ibérica, Revista de la Asociación Europea de Lenguas para Fines Específicos* 16 (2008), pp. 129–145.
- [30] Susan Šarčević. “Multilingual lawmaking and legal (un) certainty in the European Union”. In: *International journal of law, language & discourse* 3.1 (2013), pp. 1–29.
- [31] Colin Robertson. “EU law and semiotics”. In: *International Journal for the Semiotics of Law-Revue internationale de Sémiotique juridique* 23.2 (2010), pp. 145–164.

- [32] Commission of the European Communities. *Multilingualism: An Asset for Europe and a Shared Commitment*. Communication COM (2008) 566 final. Commission of the European Communities, 2008. URL: <https://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=COM:2008:0566:FIN:EN:PDF>.
- [33] Cornelis Jaap Wouter Baaij et al. *Legal integration and language diversity: the case for source-oriented EU translation*. Oxford University Press, 2018.
- [34] Alain Fenet. “Diversité linguistique et construction européenne”. In: *Revue trimestrielle de droit européen* 37.2 (2001), pp. 235–270.
- [35] Deirdre Curtin and Päivi Leino-Sandberg. *Openness, Transparency and the Right of Access to Documents in the EU. In-depth analysis for the PETI Committee*. Research paper PE 556.973. Brussels: European Parliament, 2016. URL: [https://www.europarl.europa.eu/RegData/etudes/IDAN/2016/556973/IPOL_IDA\(2016\)556973_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/IDAN/2016/556973/IPOL_IDA(2016)556973_EN.pdf).
- [36] Antoni Milian i Massana. “Le principe d’égalité des langues au sein des institutions de l’Union européenne et dans le droit communautaire, mythe ou réalité?” In: *Revista de llengua i dret* 38 (2002), pp. 47–94.
- [37] European Commission. *European Governance — A White Paper*. White Paper COM(2001) 428 final. Commission of the European Communities, July 25, 2001. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52001DC0428>.
- [38] Lucía Pacho Aljanati. *The Court of Justice of the European Union’s case law on linguistic divergences (2007-2013): interpretation criteria and implications for the translation of EU legislation*. Universitat Autònoma de Barcelona, 2015.
- [39] *Pubblico Ministero v Ratti*. Apr. 5, 1979. URL: <https://curia.europa.eu/juris/showPdf.jsf?text=&docid=88252&pageIndex=0&doclang=EN&mode=req&dir=&occ=first&part=1&cid=123456>.

- [40] *Judgment of the Court (Grand Chamber) of 11 December 2007 - Skoma-Lux sro v Celní ředitelství Olomouc*. Dec. 11, 2007. URL: <https://curia.europa.eu/juris/showPdf.jsf?text=&docid=67610&pageIndex=0&doclang=EN&mode=req&dir=&occ=first&part=1&cid=123456>.
- [41] Phoebus Athanassiou. “The application of multilingualism in the European Union context”. In: (2006).
- [42] Silvia Ferreri. “14 Multilingual interpretation of European Union law”. In: *Handbook of communication in the legal sphere* 14 (2018), p. 373.
- [43] Andrea Dahmen et al. *Translation at the European Commission - A History*. History Report. Available online: [EU Publications](https://op.europa.eu/detail/-/publication/276bad1c-5126-11ec-b7c3-01aa75ed71a1/language-en). Brussels, Belgium: Directorate-General for Translation, European Commission, 2021.
- [44] European Parliament. *Framework Strategy for Multilingualism*. European Parliament resolution on a new framework strategy for multilingualism (2006/2083(INI)). Nov. 2006. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52006IP0488>.
- [45] Bruno De Witte. “Language Law of the European Union: Protecting or Eroding Linguistic Diversity?” In: (2004).
- [46] Richard L Creech. *Law and language in the European Union: the paradox of a Babel” united in diversity*. Europa Law Publishing, 2005.
- [47] Stefaan Van der Jeught. *EU language law*. Europa Law Publishing, 2015.
- [48] Jacques Rideau. “L’Union européenne face aux défis linguistiques”. In: *SUEUR, JJ (Zuz.), Interpréter&Traduire. Actes du colloque international des*. Vol. 25. 2007, pp. 63–137.
- [49] Gálvez Salvador and María José. “La diversidad lingüística comunitaria: aspectos constitucionales”. In: (2011).

- [50] Mattias Derlén. “In defence of (limited) multilingualism: problems and possibilities of the multilingual interpretation of European Union law in national courts”. In: *Linguistic diversity and European democracy*. Routledge, 2016, pp. 143–166.
- [51] Karin Luttermann. “Multilingualism in the European Union Status quo and perspectives: The reference language model”. In: *Grewendorf (Ed.), Formal Linguistics and Law* (2009), pp. 315–338.
- [52] Rafał Mańko. *Legal aspects of EU multilingualism*. Briefing PE 595.914. European Parliamentary Research Service, Jan. 2017. URL: [https://www.europarl.europa.eu/RegData/etudes/BRIE/2017/595914/EPRS_BRI\(2017\)595914_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/BRIE/2017/595914/EPRS_BRI(2017)595914_EN.pdf).
- [53] Emma Wagner, Svend Bech, and Jesús Martínez. *Translating for the European Union institutions*. Routledge, 2014.
- [54] Normand Labrie. “La construction linguistique de la Communauté européenne”. In: *(No Title)* (1993).
- [55] Elena Ioriatti et al. *Interpretazione comparante e multilinguismo europeo*. Cedam, 2013.
- [56] J-C Piris. “Union européenne: comment rédiger une législation de qualité dans 20 langues et pour 25 États membres”. In: *Revue du droit public et de la science politique en France et a l’etranger* 2 (2005), pp. 475–492.
- [57] Tito Gallas. “Understanding EC law as ‘Diplomatic law’ and its language”. In: *Multilingualism and the Harmonisation of European Law Kluwer Law International* (2006), pp. 119–128.
- [58] Judit Várnai, Anna Bérczi, and Réka Somssich. *Lawmaking in the EU Multilingual Environment*. Study HC-31-10-678-EN-C. Available online: [EU Publications](https://op.europa.eu/en/publication-detail/-/publication/276bad1c-5126-11ec-b7c3-01aa75ed71a1/language-en). Luxembourg: Directorate-General for Translation, European Commission, Oct. 2010.

- [59] Lawrence M Solan. “Multilingualism and morality in statutory interpretation”. In: *Language & Law/Linguagem e Direito* 1.1 (2014).
- [60] A. Felici. “Translating EU Law: A New Perspective to the Paradox of Multilingualism”. In: *Comparative Legilinguistics* 2 (2010). Available online, pp. 153–165.
- [61] Robert Schütze. *European Union Law*. Oxford University Press, 2021.
- [62] Wim Voermans, Josephine MR Hartmann, and Michael Kaeding. “The quest for legitimacy in EU secondary legislation”. In: *The Theory and Practice of Legislation* 2.1 (2014), pp. 5–32.
- [63] European Parliament. *Overview of the Ordinary Legislative Procedure*. Web page. 2023. URL: <https://www.europarl.europa.eu/olp/en/ordinary-legislative-procedure/overview>.
- [64] Council of the European Union. *Joint Declaration on Practical Arrangements for the Codecision Procedure (Article 251 of the EC Treaty)*. Official Journal of the European Union. June 2007. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32007C0630%2801%29>.
- [65] Gijs Jan Brandsma. “Transparency of EU informal trilogues through public feedback in the European Parliament: Promise unfulfilled”. In: *Journal of European Public Policy* 26.10 (2019), pp. 1464–1483.
- [66] Anna Doliwa-Klepcka. “Preparation of a legislative proposal in the multilingual legal system of the European Union”. In: *Studies in Logic, Grammar and Rhetoric* 45.1 (2016), pp. 37–50.
- [67] William Robinson. “How the European Commission drafts legislation in 20 languages”. In: *Clarity—Journal of the International Association Promoting Plain Legal Language* 53 (2005), pp. 4–10.
- [68] William Robinson. *Drafting European Union Legislation*. In-Depth Analysis PE 462.442. Requested by the European Parliament’s Legal Affairs Committee. Brussels: European Parliament, Policy Department C: Citizens’ Rights and Constitutional Affairs, Apr. 2012.

- [69] Michele Graziadei. “Many Languages for a Single Voice”. In: *Drafting the CESL* (2014).
- [70] Stefania Dragone. “The quality of Community legislation and the role of the European Commission legal revisers”. In: *Multilingualism and the Harmonization of European Law. Alphen aan den Rijn: Kluwer Law International* (2006).
- [71] Lara Panning. “Building and managing the European Commission’s position for trilogue negotiations”. In: *Journal of European Public Policy* 28.1 (2021), pp. 32–52.
- [72] William Robinson. “Drafting EU legislation in the European Commission: A collaborative process”. In: *The Theory and Practice of Legislation* 2.3 (2014), pp. 249–272.
- [73] Manuela Guggeis. “How and when lawyer-linguists of the EU institutions intervene during the legislative procedure for the adoption of the regulation on a Common European Sales Law (CESL)”. In: *Drafting the CESL* (2014).
- [74] Manuela Guggeis. “Multilingualism in the European Union Decision-Making Process”. In: *Criminal Proceedings, Languages and the European Union: Linguistic and Legal Issues*. Springer, 2013, pp. 45–56.
- [75] Jaroslav G. Polach. “Harmonization of Laws in Western Europe”. In: *The American Journal of Comparative Law* 8.2 (1959), pp. 148–167. ISSN: 0002919X. URL: <http://www.jstor.org/stable/837456> (visited on 05/31/2024).
- [76] Elena Ruiz-Cortés. “Legal-linguistic analysis of EU law and its transpositions: A useful approach to explore harmonisation problems?” In: *Journal of Language & Law/Revista de Llengua i Dret* 74 (2020).
- [77] Robert Thomas. “Transposing European Union Law in the United Kingdom: Administrative Rule-Making, Scrutiny and Better Regulation”. In: *European Public Law* 14.2 (2008).

- [78] Colin Robertson. “How the European Union functions in 23 languages”. In: (2013).
- [79] *Cross-Justice: Legal Studies and Research*. <https://site.unibo.it/cross-justice/en>. Accessed: 2024-09-10. 2024.
- [80] *Facilex: The Project*. <https://site.unibo.it/facilex/en/the-project>. Accessed: 2024-09-10. 2024.
- [81] Alexander Roßnagel et al. “National implementation of the General Data Protection Regulation: challenges-approaches-strategies: policy paper”. In: *National implementation of the General Data Protection Regulation: challenges-approaches-strategies: policy paper: Roßnagel, Alexander— uBile, Tamer— uFriedewald, Michael— uGeminn, Christian— uGrigorjew, Olga— uKaraboga, Murat— uNebel, Mari*. Karlsruhe: Fraunhofer-Institut für System-und Innovationsforschung ISI, 2018.
- [82] Deborah Cao. “Inter-lingual uncertainty in bilingual and multilingual law”. In: *Journal of Pragmatics* 39.1 (2007), pp. 69–83.
- [83] Elina Paunio and Susanna Lindroos-Hovinheimo. “Taking language seriously: An analysis of linguistic reasoning and its implications in EU law”. In: *European Law Journal* 16.4 (2010), pp. 395–416.
- [84] Ingemar Strandvik. “Evaluation of outsourced translations. State of play in the European Commission’s Directorate-General for Translation (DGT)”. In: *Quality aspects in institutional translation* 8 (2017), pp. 123–137.
- [85] Ingemar Strandvik. “Is there scope for a more professional approach to EU multilingual lawmaking?” In: *The theory and practice of legislation* 2.2 (2014), pp. 211–228.
- [86] Directorate-General for Translation European Commission. *Quantifying Quality Costs and the Cost of Poor Quality in Translation*. Study on Translation and Multilingualism HC-31-12-463-EN-C. Available online: <https://op.europa.eu/en/publication-detail/->

/publication/276bad1c-5126-11ec-b7c3-01aa75ed71a1/language-en.
Luxembourg: European Commission, 2012.

- [87] Lucja Biel and Jan Engberg. “Research models and methods in legal translation”. In: *Linguistica Antverpiensia, New Series—Themes in Translation Studies* 12 (2013).
- [88] Anne Lise Kjær. “Legal Translation in the European Union: A Research Field in Need of a New Approach”. English. In: *Language and the Law*. Ed. by Krzysztof Kredens and Stanislaw Gozdz-Roszkowski. Lodz Studies in Language 16. Peter Lang, 2007, pp. 69–95. ISBN: 9783631574478.
- [89] Elina Paunio. *Legal certainty in multilingual EU law: language, discourse and reasoning at the European Court of Justice*. Routledge, 2016.
- [90] Lawrence M Solan. *Vagueness and ambiguity in legal interpretation*. na, 2005.
- [91] Fernando Prieto Ramos. “International and supranational law in translation: From multilingual lawmaking to adjudication”. In: *The Translator* 20.3 (2014), pp. 313–331.
- [92] Theodor Schilling. “Beyond multilingualism: On different approaches to the handling of diverging language versions of a community law”. In: *European Law Journal* 16.1 (2010), pp. 47–66.
- [93] Simon Blackburn. *The Oxford dictionary of philosophy*. OUP Oxford, 2005.
- [94] Vilelmini Sosoni. “Training translators to work for the EU institutions: luxury or necessity”. In: *The Journal of Specialised Translation* 16 (2011), pp. 77–108.
- [95] Lucja Biel. “Translation of multilingual EU legislation as a sub-genre of legal translation”. In: (2007).

- [96] Giovanni Sartor et al. *The Nature of Legal Concepts: Inferential Nodes Or Ontological Categories?*. European University Institute, 2007.
- [97] Sacha Prechal. “Direct effect, indirect effect, supremacy and the evolving constitution of the European Union”. In: *The fundamentals of EU law revisited: Assessing the impact of the constitutional debate* (2007), pp. 35–71.
- [98] Gunnar Beck. *The Legal Reasoning of the Court of Justice of the EU*. Bloomsbury Publishing, 2013.
- [99] Theodore George. “Hermeneutics”. In: (2020).
- [100] Gerrit Betlem. “The principle of indirect effect of Community law”. In: *European Review of private law* 3.1 (1995).
- [101] Jerzy Wróblewski. “Semantic Basis of the Theory of Legal interpretation”. In: *Logique et analyse* 6.21/24 (1963), pp. 397–416.
- [102] Timothy Endicott. “Law and Language”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Spring 2022. Metaphysics Research Lab, Stanford University, 2022.
- [103] Ross Charnock. “Hart as contextualist? Theories of interpretation in language and the law”. In: *Law and language: Current legal issues* 15 (2013), pp. 128–150.
- [104] Richard H Fallon Jr. “The Meaning of Legal Meaning and Its Implications for Theories of Legal Interpretation”. In: *U. Chi. L. Rev.* 82 (2015), p. 1235.
- [105] John Perry. “Textualism and the Discovery of Rights”. In: (2011).
- [106] Lawrence B Solum. “Communicative content and legal content”. In: *Notre Dame L. Rev.* 89 (2013), p. 479.
- [107] Michael S Moore. “A natural law theory of interpretation”. In: *Law and Legal Interpretation*. Taylor and Francis Inc., 2017, pp. 113–234.

- [108] Ditlev Tamm. “The history of the Court of Justice of the European Union since its origin”. In: *The Court of Justice and the Construction of Europe: Analyses and Perspectives on Sixty Years of Case-law-La Cour de Justice et la Construction de l’Europe: Analyses et Perspectives de Soixante Ans de Jurisprudence* (2013), pp. 9–35.
- [109] Giulio Itzcovich. “The interpretation of community law by the European Court of Justice”. In: *German Law Journal* 10.5 (2009), pp. 537–560.
- [110] Albertina Albors-Llorens. “Judicial protection before the Court of Justice of the European Union”. In: *European Union Law* (2014), pp. 255–299.
- [111] European Union. *The Direct Effect of European Union Law*. Web page. 2023. URL: <https://eur-lex.europa.eu/IT/legal-content/summary/the-direct-effect-of-european-union-law.html>.
- [112] Court of Justice of the European Union. *Van Gend en Loos v Nederlandse Administratie der Belastingen*. Judgment of the Court of 5 February 1963. Case 26/62. 1963. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:61962CJ0026>.
- [113] Court of Justice of the European Union. *Von Colson and Kamann v Land Nordrhein-Westfalen*. Judgment of the Court of 10 April 1984. Case 14/83. 1984. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:61983CJ0014>.
- [114] Isabel Schübel-Pfister. *Sprache und Gemeinschaftsrecht*. Duncker und Humblot, 2019.
- [115] Sieg Eiselen. “Literal interpretation: the meaning of the words”. In: *en A. Janssen & O. Meyer (ed.). CISG Methodology. Munich. Sellier* (2009), pp. 61–89.
- [116] Anthony Arnall. *The European Union and its court of justice*. Oxford University Press, 2006.

- [117] Koen Lenaerts. “Rule of Law and the Coherence of the Judicial System of the European Union, The”. In: *Common Market L. Rev.* 44 (2007), p. 1625.
- [118] Court of Justice of the European Union. *Pfeiffer and Others v Deutsches Rotes Kreuz, Kreisverband Waldshut eV*. Judgment of the Court of 5 October 2004. Case C-336/03. 2004. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:62003CJ0336>.
- [119] Samuli Miettinen and Merita Kettunen. “Travaux to the EU Treaties: Preparatory Work as a Source of EU Law”. In: *Cambridge Yearbook of European Legal Studies* 17 (2015), pp. 145–167.
- [120] T. Klimas and J. Vaiciukaite. “The Law of Recitals in European Community Legislation”. In: *ILSA Journal of International & Comparative Law* 15.6 (July 2008), pp. 61–93.
- [121] Pierre Pescatore. *Les Objectifs de la Communauté européenne comme principe d’interprétation dans la jurisprudence de la Cour de Justice: Miscellanea WJ Ganshof van der Meersch. Studia ab discipulis amicisque in honorem egregii professoris edita*. E. Bruylant, 1972.
- [122] Tamás Szabados. “The Role of Language in Legal Interpretation: The Case Law of the Court of Justice of the European Union”. In: *T Szabados, The Role of Language in Legal Interpretation: The Case Law of the Court of Justice of the European Union, in R. Somssich and T. Szabados (eds), Central and Eastern European Countries after and before the Accession* 1 (2011).
- [123] Joxerramon Bengoetxea. “Multilingual and multicultural legal reasoning: the European Court of Justice”. In: *Linguistic diversity and European democracy*. Routledge, 2016, pp. 97–122.
- [124] *Judgment of the Court of 15 July 1964 - Flaminio Costa v E.N.E.L.* July 15, 1964.

- [125] Sieglinde E Pommer. “Interpreting Multilingual EU Law: What Role for Legal Translation?” In: *Eur. Rev. Private L.* 20 (2012), p. 1241.
- [126] Susanna Lindroos-Hovinheimo. “On the Indeterminacy of Legal Translation”. In: *Private law and the many cultures of Europe*. Kluwer Law International, 2007, pp. 367–383.
- [127] Manuela Guggeis and William Robinson. “Co-Revision”: Legal-Linguistic Revision in the European Union “Co-Decision” Process”. In: *The role of legal translation in legal harmonization* (2012), pp. 51–81.
- [128] Olga Łachacz and Rafał Mańko. “Multilingualism at the Court of Justice of the European Union: theoretical and practical aspects”. In: *Studies in Logic, Grammar and Rhetoric* 34.1 (2013), pp. 75–92.
- [129] Randy J Kozel. “Stare decisis as judicial doctrine”. In: *Wash. & Lee L. Rev.* 67 (2010), p. 411.
- [130] Duncan Robertson. “Agent-Based Modeling Toolkits NetLogo, RePast, and Swarm”. In: *Academy of Management Learning and Education, The* 4 (Dec. 2005), pp. 524–527. DOI: [10.5465/AMLE.2005.19086798](https://doi.org/10.5465/AMLE.2005.19086798).
- [131] European Union. *IATE - Interactive Terminology for Europe*. Web page. Accessed on [your access date]. 2023. URL: <https://iate.europa.eu/home>.
- [132] Publications Office of the European Union. *Cellar Data*. Web page. 2023. URL: <https://op.europa.eu/en/web/cellar/cellar-data>.
- [133] Publications Office of the European Union. *SPARQL Cellar of the Publications Office*. Web page. 2023. URL: <https://data.europa.eu/data/datasets/sparql-cellar-of-the-publications-office?locale=en>.

- [134] Publications Office of the European Union. *Dataset of Akoma Ntoso for the European Union (AKN4EU)*. Web page. 2023. URL: <https://op.europa.eu/it/web/eu-vocabularies/dataset/-/resource?uri=http://publications.europa.eu/resource/dataset/akn4eu>.
- [135] General Secretariat of the Council of the European Union. *The Language Service of the General Secretariat of the Council of the European Union. Making Multilingualism Work*. Language Service Report QC-32-11-696-EN. Available online: <https://www.consilium.europa.eu>. Rue de la Loi/Wetstraat 175, 1048 Bruxelles/Brussel, Belgique/België: General Secretariat of the Council of the European Union, Feb. 2012. DOI: [10.2860/79193](https://doi.org/10.2860/79193).
- [136] José Manuel Barroso, Domenico Cosmai, and David Albert Best. *The Language of Europe. Multilingualism and translation in the EU institutions: Practice, problems and perspectives*. Editions de l’Université de Bruxelles, 2014.
- [137] European Commission. *eTranslation - Multilingual Machine Translation for European Union Public Services and Institutions*. Web page. 2023. URL: https://commission.europa.eu/resources-partners/etranslation_en.
- [138] Nicholas Pipitone and Ghita Houir Alami. “LegalBench-RAG: A Benchmark for Retrieval-Augmented Generation in the Legal Domain”. In: *arXiv preprint arXiv:2408.10343* (2024).
- [139] Haoxi Zhong et al. “How Does NLP Benefit Legal System: A Summary of Legal Artificial Intelligence”. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*. Ed. by Dan Jurafsky et al. Association for Computational Linguistics, 2020, pp. 5218–5230. DOI: [10.18653/v1/2020.acl-main.466](https://doi.org/10.18653/v1/2020.acl-main.466).
- [140] Roland Friedrich, Mauro Luzzatto, and Elliott Ash. “Entropy in Legal Language”. In: *Proceedings of the Natural Legal Language*

Processing Workshop 2020 co-located with the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD 2020), Virtual Workshop, August 24, 2020. Ed. by Nikolaos Aletras et al. Vol. 2645. CEUR Workshop Proceedings. CEUR-WS.org, 2020, pp. 25–30. URL: <http://ceur-ws.org/Vol-2645/paper4.pdf>.

- [141] Adam Wyner et al. “Approaches to text mining arguments from legal cases”. In: *Semantic processing of legal texts*. Springer, 2010, pp. 60–79.
- [142] Guido Boella et al. “A system for classifying multi-label text into EuroVoc”. In: *Proceedings of the fourteenth international conference on artificial intelligence and law*. 2013, pp. 239–240.
- [143] Guido Boella et al. “Eunomos, a legal document and knowledge management system for the web to provide relevant, reliable and up-to-date information on the law”. In: *Artificial Intelligence and Law* 24.3 (2016), pp. 245–283.
- [144] Adebayo Kolawole John et al. “Legalbot: A deep learning-based conversational agent in the legal domain”. In: *International Conference on Applications of Natural Language to Information Systems*. Springer. 2017, pp. 267–273.
- [145] Livio Robaldo et al. *Introduction for artificial intelligence and law: special issue “natural language processing for legal texts”*. 2019. DOI: [doi10.1007/s10506-019-09251-2](https://doi.org/10.1007/s10506-019-09251-2).
- [146] Guido Boella, Luigi Di Caro, and Valentina Leone. “Semi-automatic knowledge population in a legal document management system”. In: *Artificial Intelligence and Law* 27.2 (2019), pp. 227–251.
- [147] Nirmalie Wiratunga et al. “CBR-RAG: case-based reasoning for retrieval augmented generation in LLMs for legal question answering”. In: *International Conference on Case-Based Reasoning*. Springer. 2024, pp. 445–460.

- [148] Davide Audrito and Andrea Filippo Ferraris. “Legal Design through Business Process Model and Notation (BPMN): a Digital Services Act case-study.” In: *JOWO*. 2023.
- [149] Ashish Chouhan and Michael Gertz. “LexDrafter: Terminology Drafting for Legislative Documents using Retrieval Augmented Generation”. In: *arXiv preprint arXiv:2403.16295* (2024).
- [150] Rosa Meo and Emilio Sulis. “Processing Affect in Social Media: A Comparison of Methods to Distinguish Emotions in Tweets”. In: *ACM Trans. Internet Techn.* 17.1 (2017), 7:1–7:25. DOI: [10.1145/2996187](https://doi.org/10.1145/2996187). URL: <https://doi.org/10.1145/2996187>.
- [151] Emilio Sulis et al. “Exploring Sentiment in Social Media and Official Statistics: a General Framework”. In: *Proceedings of the 2nd International Workshop on Emotion and Sentiment in Social and Expressive Media: Opportunities and Challenges for Emotion-aware Multiagent Systems co-located with 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Istanbul, Turkey, May 5, 2015*. Ed. by Cristina Bosco et al. Vol. 1351. CEUR Workshop Proceedings. CEUR-WS.org, 2015, pp. 96–105. URL: <http://ceur-ws.org/Vol-1351/paper8.pdf>.
- [152] Ronen Feldman and James Sanger. *The Text Mining Handbook - Advanced Approaches in Analyzing Unstructured Data*. Cambridge University Press, 2007. ISBN: 978-0-521-83657-9.
- [153] Daniel Martin Katz, Ron Dolin, and Michael J Bommarito. *Legal Informatics*. Cambridge University Press, 2021.
- [154] Massimo Durante. *Computational Power: The Impact of ICT on Law, Society and Knowledge*. Routledge, 2021.
- [155] Mariangela Biasiotti et al. “Legal informatics and management of legislative documents”. In: *Global Center for ICT in Parliament Working Paper 2* (2008).

- [156] Santiago Ontañón. “An overview of distance and similarity functions for structured data”. In: *Artificial Intelligence Review* 53.7 (2020), pp. 5309–5351. DOI: <https://doi.org/10.1007/s10462-020-09821-w>.
- [157] Mi-Young Kim, Ying Xu, and Randy Goebel. “Legal question answering using ranking svm and syntactic/semantic similarity”. In: *JSAI International Symposium on Artificial Intelligence*. Springer. 2014, pp. 244–258.
- [158] Rupali Wagh and Deepa Anand. “Application of citation network analysis for improved similarity index estimation of legal case documents : A study”. In: *2017 IEEE International Conference on Current Trends in Advanced Computing (ICCTAC)*. 2017, pp. 1–5. DOI: [10.1109/ICCTAC.2017.8249996](https://doi.org/10.1109/ICCTAC.2017.8249996).
- [159] Markus Haverland, Bernard Steunenberg, and Frans Van Waarden. “Sectors at different speeds: analysing transposition deficits in the European Union”. In: *JCMS: Journal of Common Market Studies* 49.2 (2011), pp. 265–291.
- [160] Llio Humphreys et al. “Mapping Recitals to Normative Provisions in EU Legislation to Assist Legal Interpretation.” In: *JURIX*. 2015, pp. 41–49.
- [161] R. Nanda et al. “Unsupervised and supervised text similarity systems for automated identification of national implementing measures of European directives”. In: *Artificial Intelligence and Law* 27 (2018), pp. 199–225. DOI: <https://doi.org/10.1007/s10506-018-9236-y>.
- [162] Omer Levy, Yoav Goldberg, and Ido Dagan. “Improving Distributional Similarity with Lessons Learned from Word Embeddings”. In: *Trans. Assoc. Comput. Linguistics* 3 (2015), pp. 211–225. DOI: [10.1162/tacl_a_00134](https://doi.org/10.1162/tacl_a_00134).
- [163] Sara Renjit and Sumam Mary Idicula. “CUSAT NLP@AILA-FIRE2019: Similarity in Legal Texts using Document Level Embeddings”. In:

- Working Notes of FIRE 2019 - Forum for Information Retrieval Evaluation, Kolkata, India, December 12-15, 2019*. Ed. by Parth Mehta et al. Vol. 2517. CEUR Workshop Proceedings. CEUR-WS.org, 2019, pp. 25–30. URL: <http://ceur-ws.org/Vol-2517/T1-4.pdf>.
- [164] Rupali S. Wagh and Deepa Anand. “Legal document similarity: a multi-criteria decision-making perspective”. In: *PeerJ Comput. Sci.* 6 (2020), e262. DOI: [10.7717/peerj-cs.262](https://doi.org/10.7717/peerj-cs.262). URL: <https://doi.org/10.7717/peerj-cs.262>.
- [165] Sushanta Kumar et al. “Finding Similar Legal Judgements under Common Law System”. In: *Databases in Networked Information Systems - 8th International Workshop, DNIS 2013, Aizu-Wakamatsu, Japan, March 25-27, 2013. Proceedings*. Ed. by Aastha Madaan, Shinji Kikuchi, and Subhash Bhalla. Vol. 7813. Lecture Notes in Computer Science. Springer, 2013, pp. 103–116. DOI: [10.1007/978-3-642-37134-9_9](https://doi.org/10.1007/978-3-642-37134-9_9). URL: https://doi.org/10.1007/978-3-642-37134-9_9.
- [166] Emilio Sulis et al. “Exploiting co-occurrence networks for classification of implicit inter-relationships in legal texts”. In: *Information Systems* (2021), p. 101821. ISSN: 0306-4379. DOI: <https://doi.org/10.1016/j.is.2021.101821>.
- [167] Arpan Mandal et al. “Unsupervised approaches for measuring textual similarity between legal court case reports”. In: *Artif. Intell. Law* 29.3 (2021), pp. 417–451. DOI: [10.1007/s10506-020-09280-2](https://doi.org/10.1007/s10506-020-09280-2).
- [168] Stefan Vogenauer and Stephen Weatherill. *The harmonisation of European contract law: implications for European private laws, business and legal practice*. Bloomsbury Publishing, 2006. DOI: <https://doi.org/10.1111/j.1468-0386.2007.00376.x>.
- [169] Mads Andenas and Camilla Baasch Andersen. *Theory and practice of harmonisation*. Edward Elgar Publishing, 2012.

- [170] Helmut Satzger. “The Harmonisation of Criminal Sanctions in the European Union - A New Approach”. In: *eu crim* (2019). DOI: [10.30709/eu crim-2019-007](https://doi.org/10.30709/eu crim-2019-007). URL: <https://doi.org/10.30709/eu crim-2019-007>.
- [171] Werner Schroeder. “Limits to European Harmonisation of Criminal Law”. In: *eu crim* (2020). DOI: [10.30709/eu crim-2020-008](https://doi.org/10.30709/eu crim-2020-008). URL: <https://doi.org/10.30709/eu crim-2020-008>.
- [172] Giulia Lasagni et al. “The Italian Implementation of the EU Directives on Procedural Safeguards for Accused Persons in Criminal Proceedings”. In: *SISTEMA PENALE* (2022), pp. 1–139.
- [173] Christian Kaunert, John D Occhipinti, and Sarah Léonard. *Introduction: supranational governance in the Area of Freedom, Security and Justice after the Stockholm Programme*. 2014.
- [174] B. Steunenbergh and Mark Rhinard. “The transposition of European law in EU member states: between process and politics”. In: *European Political Science Review* 2 (2010), pp. 495–520. DOI: <https://doi.org/10.1017/S1755773910000196>.
- [175] Dionyssis G Dimitrakopoulos. “The transposition of EU law: ‘post-decisional politics’ and institutional autonomy”. In: *European Law Journal* 7.4 (2001), pp. 442–458.
- [176] Cornelis J Van Rijsbergen, Stephen Edward Robertson, and Martin F Porter. *New models in probabilistic information retrieval*. Vol. 5587. British Library Research and Development Department London, 1980.
- [177] Michael AA Cox and Trevor F Cox. “Multidimensional scaling”. In: *Handbook of data visualization*. Springer, 2008, pp. 315–347.
- [178] Emilio Sulis et al. “Exploring Network Analysis in a Corpus-Based Approach to Legal Texts: A Case Study”. In: *Proceedings of the First International Workshop “CAiSE for Legal Documents” (COUrT 2020) co-located with the 32nd International Conference on Advanced Information Systems Engineering (CAiSE 2020), Grenoble,*

- France, June 9, 2020. Ed. by Andrea Tagarelli et al. Vol. 2690. CEUR Workshop Proceedings. CEUR-WS.org, 2020, pp. 27–38. URL: <http://ceur-ws.org/Vol-2690/C0UrT-paper3.pdf>.
- [179] Ilaria Angela Amantea et al. “Modelling Norm Types and their Inter-relationships in EU Directives”. In: *Proceedings of the Third Workshop on Automated Semantic Analysis of Information in Legal Texts co-located with the 17th International Conference on Artificial Intelligence and Law (ICAAIL 2019), Montreal, QC, Canada, June 21, 2019*. Ed. by Kevin D. Ashley et al. Vol. 2385. CEUR Workshop Proceedings. CEUR-WS.org, 2019. URL: <http://ceur-ws.org/Vol-2385/paper8.pdf>.
- [180] Thomas F Gordon. “The Role of Exceptions in Models of the Law”. In: *Formalisierung im Recht und Ansätze juristischer Expertensysteme* (1986), pp. 52–59.
- [181] Mark Sainsbury. *Logical forms: An introduction to philosophical logic*. John Wiley & Sons, 2000.
- [182] Katrin Tomanek, Joachim Wermter, and Udo Hahn. “An approach to text corpus construction which cuts annotation costs and maintains reusability of annotated data”. In: *Proceedings of the 2007 EMNLP-CoNLL*. 2007, pp. 486–495.
- [183] L Karl Branting. “Data-centric and logic-based models for automated legal problem solving”. In: *Artificial Intelligence and Law* 25.1 (2017), pp. 5–27.
- [184] Henry Prakken. “A logical framework for modelling legal argument”. In: *Proceedings of the 4th international conference on Artificial intelligence and law*. 1993, pp. 1–9.
- [185] Guido Boella et al. “Eunomos, a Legal Document and Knowledge Management System for the Web to Provide Relevant, Reliable and up-to-Date Information on the Law”. In: *Artificial Intelligence and Law* 24.3 (Sept. 2016), pp. 245–283. ISSN: 0924-8463.

- [186] Rada Mihalcea and Dragomir Radev. *Graph-based natural language processing and information retrieval*. Cambridge university press, 2011.
- [187] Paul Zhang and Lavanya Koppaka. “Semantics-Based Legal Citation Network”. In: *Proceedings of the 11th International Conference on Artificial Intelligence and Law*. ICAIL ’07. Stanford, California: Association for Computing Machinery, 2007, pp. 123–130. ISBN: 9781595936806.
- [188] Rohan Nanda et al. “Multilingual Legal Information Retrieval System for Mapping Recitals and Normative Provisions”. In: *JURI SAYS* (2020), p. 123.
- [189] Ali Sadeghian et al. “Semantic edge labeling over legal citation graphs”. In: *Proceedings of the workshop on legal text, document, and corpus analytics (LTDCA-2016)*. 2016, pp. 70–75.
- [190] Marios Koniaris, Ioannis Anagnostopoulos, and Yannis Vassiliou. “Network Analysis in the Legal Domain: A complex model for European Union legal sources”. In: *Journal of Complex Networks* 6.2 (2018), pp. 243–268.
- [191] Maarten den Heijer, Teun van Os van den Abeelen, and Antanina Maslyka. “On the Use and Misuse of Recitals in European Union Law”. In: *SSRN Electronic Journal* 1.2019-31 (2019).
- [192] Koen Lenaerts. “Interpretation and the Court of Justice: A Basis for Comparative Reflection”. In: *Int’l Law*. 41 (2007), p. 1011.
- [193] *Judgment of the Court (Fifth Chamber) of 19 November 1998 - Criminal proceedings against Gunnar Nilsson, Per Olov Hagelgren and Solweig Arrborn*. Case C-162/97. European Court Reports 1998, Page I-07473. 1998.
- [194] *Judgment of the Court (Grand Chamber) of 10 January 2006 - The Queen, on the application of International Air Transport Association and European Low Fares Airline Association v Department*

for Transport. Case C-344/04. European Court Reports 2006, Page I-00403. 2006.

- [195] *Judgment of the Court of 4 July 1963 - Federal Republic of Germany v Commission of the European Economic Community*. Case 24/62. European Court Reports 1963, Page 00063. 1963.
- [196] Sylvia Mercado Kierkegaard. “Here comes the ‘cybernators’!” In: *Computer Law & Security Review* 22.5 (2006), pp. 381–391.
- [197] Jack G Conrad and L Karl Branting. “Introduction to the special issue on legal text analytics”. In: *Artificial Intelligence and Law* 26.2 (2018), pp. 99–102.
- [198] Kevin D Ashley. *Artificial intelligence and legal analytics: new tools for law practice in the digital age*. Cambridge University Press, 2017.
- [199] Mireille Hildebrandt. “The meaning and the mining of legal texts”. In: *Understanding Digital Humanities*. Springer, 2012, pp. 145–160.
- [200] Panteleimon Krasadakis, Evangelos Sakkopoulos, and Vassilios S Verykios. “A survey on challenges and advances in natural language processing with a focus on legal informatics and low-resource languages”. In: *Electronics* 13.3 (2024), p. 648.
- [201] Ingo Glaser, Elena Scepankova, and Florian Matthes. “Classifying Semantic Types of Legal Sentences: Portability of Machine Learning Models.” In: *JURIX*. 2018, pp. 61–70.
- [202] Jerrold Soh, How Khang Lim, and Ian Ernst Chai. “Legal Area Classification: A Comparative Study of Text Classifiers on Singapore Supreme Court Judgments”. In: *Proceedings of the Natural Legal Language Processing Workshop 2019*. Minneapolis, Minnesota: Association for Computational Linguistics, June 2019, pp. 67–77.
- [203] Bernhard Walzl et al. “Classifying Legal Norms with Active Machine Learning.” In: *JURIX*. 2017, pp. 11–20.

- [204] Bernhard Walzl et al. “Semantic types of legal norms in German laws: classification and analysis using local linear explanations”. In: *Artificial Intelligence and Law* 27.1 (2019), pp. 43–71.
- [205] T. T. Cheng et al. “Information extraction from legal documents”. In: *2009 Eighth International Symposium on Natural Language Processing*. 2009, pp. 157–162.
- [206] Ilias Chalkidis and Dimitrios Kampas. “Deep learning in law: early adaptation and legal word embeddings trained on large corpora”. In: *Artificial Intelligence and Law* 27.2 (2019), pp. 171–198.
- [207] Arpan Mandal et al. “Unsupervised approaches for measuring textual similarity between legal court case reports”. In: *Artificial Intelligence and Law* (2021), pp. 1–35.
- [208] Daniel Martin Katz, Michael J Bommarito II, and Josh Blackman. “Predicting the behavior of the supreme court of the united states: A general approach”. In: *arXiv preprint arXiv:1407.6333* (2014).
- [209] Masha Medvedeva, Michel Vols, and Martijn Wieling. “Using machine learning to predict decisions of the European Court of Human Rights”. In: *Artificial Intelligence and Law* 28.2 (2020), pp. 237–266.
- [210] Damian Langone, Alfredo Fulloni, and Dina Wonsever. “A citations network for legal decisions”. Jurisin conference. 2020. URL: <http://research.nii.ac.jp/~ksatoh/jurisin2020/>.
- [211] Daniele Bourcier and Pierre Mazzega. “Codification, law article and graphs”. In: *Frontiers in Artificial Intelligence and Applications* 165 (2007), p. 29.
- [212] Morayo Adedjouma, Mehrdad Sabetzadeh, and Lionel C Briand. “Automated detection and resolution of legal cross references: Approach and a study of luxembourg’s legislation”. In: *2014 IEEE 22nd International Requirements Engineering Conference (RE)*. IEEE. 2014, pp. 63–72.

- [213] Ali Sadeghian et al. “Automatic semantic edge labeling over legal citation graphs”. In: *Artificial Intelligence and Law* 26.2 (2018), pp. 127–144.
- [214] Monica Palmirani, Raffaella Brighi, and Matteo Massini. “Automated extraction of normative references in legal texts”. In: *Proceedings of the 9th international conference on Artificial intelligence and law*. 2003, pp. 105–106.
- [215] Elena Leitner, Georg Rehm, and Julian Moreno-Schneider. “Fine-grained named entity recognition in legal documents”. In: *International Conference on Semantic Systems*. Springer. 2019, pp. 272–287.
- [216] Ngo Xuan Bach et al. “Reference extraction from Vietnamese legal documents”. In: *Proceedings of the Tenth International Symposium on Information and Communication Technology*. 2019, pp. 486–493.
- [217] Neda Sakhaee and Mark C Wilson. “Information extraction framework to build legislation network”. In: *Artificial Intelligence and Law* 29.1 (2021), pp. 35–58.
- [218] Romain Boulet, Pierre Mazzega, and Danièle Bourcier. “Network approach to the French system of legal codes part II: the role of the weights in a network”. In: *Artificial Intelligence and Law* 26.1 (2018), pp. 23–47.
- [219] Mattias Derlén and Johan Lindholm. “Goodbye van G end en L oos, Hello B osman? Using Network Analysis to Measure the Importance of Individual CJEU Judgments”. In: *European Law Journal* 20.5 (2014), pp. 667–687.
- [220] Marc van Opijnen. “Citation Analysis and Beyond: in Search of Indicators Measuring Case Law Importance.” In: *JURIX*. Vol. 250. 2012, pp. 95–104.
- [221] Marc van Opijnen. “A model for automated rating of case law”. In: *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law*. 2013, pp. 140–149.

- [222] Pierre Mazzega, Danièle Bourcier, and Romain Boulet. “The network of French legal codes”. In: *Proceedings of the 12th international conference on artificial intelligence and law*. 2009, pp. 236–237.
- [223] Guiraudé Lame. “Using NLP techniques to identify legal ontology components: Concepts and relations”. In: *Artificial Intelligence and Law* 12.4 (2004), pp. 379–396.
- [224] Jeremy C. Maxwell et al. “A Legal Cross-References Taxonomy for Reasoning about Compliance Requirements”. In: *Requir. Eng.* 17.2 (June 2012), pp. 99–115. ISSN: 0947-3602.
- [225] Yannis Panagis, Urska Sadl, and Fabien Tarissan. “Giving every case its (legal) due the contribution of citation networks and text similarity techniques to legal studies of european union law”. In: *30th International Conference on Legal Knowledge and Information Systems (JURIX’17)*. Vol. 302. IOS Press. 2017, pp. 59–68.
- [226] James H. Fowler et al. “Network Analysis and the Law: Measuring the Legal Importance of Precedents at the U.S. Supreme Court”. In: *Political Analysis* 15.3 (2007), pp. 324–346. DOI: [10.1093/pan/mpm011](https://doi.org/10.1093/pan/mpm011).
- [227] Radboud Winkels and Jelle de Ruyter. “Survival of the Fittest: Network Analysis of Dutch Supreme Court Cases”. In: *AI Approaches to the Complexity of Legal Systems. Models and Ethical Challenges for Legal Systems, Legal Language and Legal Ontologies, Argumentation and Software Agents*. Ed. by Monica Palmirani et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 106–115.
- [228] Mattias Derlén and Johan Lindholm. “Is it Good Law? Network Analysis and the CJEU’s Internal Market Jurisprudence”. In: *Journal of International Economic Law* 20.2 (May 2017), pp. 257–277.
- [229] Henrik Palmer Olsen and Aysel Küçüksu. “Finding hidden patterns in ECtHR’s case law: On how citation network analysis can improve

- our knowledge of ECtHR’s Article 14 practice”. In: *International Journal of Discrimination and the Law* 17.1 (2017), pp. 4–22.
- [230] Damian Langone, Alfredo Fulloni, and Dina Wonsever. “A citations network for legal decisions”. In: *New Frontiers in Artificial Intelligence*. Lecture Notes in Computer Science. Springer International Publishing, 2020.
 - [231] Mohammad Hamdaqa and Abdelwahab Hamou-Lhadj. “An Approach Based on Citation Analysis to Support Effective Handling of Regulatory Compliance”. In: *Future Generation Computer System* 27.4 (2011). ISSN: 0167-739X. DOI: [10.1016/j.future.2010.09.007](https://doi.org/10.1016/j.future.2010.09.007).
 - [232] Emile De Maat, Radboud Winkels, and Tom Van Engers. “Automated Detection of Reference”. In: *Legal Knowledge and Information Systems: JURIX 2006: the Nineteenth Annual Conference*. Vol. 152. IOS Press. 2006, p. 41.
 - [233] Urska Sadl and Fabien Tarissan. “The Relevance of the Network Approach to European Case Law. Reflexion and Evidence”. In: *New Legal Approaches to Studying the Court of Justice*. Dec. 2020. URL: <https://hal.archives-ouvertes.fr/hal-03098351>.
 - [234] Urška Šadl and Henrik Palmer Olsen. “Can quantitative methods complement doctrinal legal studies? Using citation network and corpus linguistic analysis to understand international courts”. In: *Leiden Journal of International Law* 30.2 (2017), pp. 327–349.
 - [235] Urska Sadl. “The role of effet utile in preserving the continuity and authority of European Union law: evidence from the citation web of the pre-accession case law of the Court of Justice of the EU”. In: *Eur. J. Legal Stud.* 8 (2015), p. 18.
 - [236] Fabien Tarissan, Yannis Panagis, and Urška Šadl. “Selecting the cases that defined Europe: complementary metrics for a network analysis”. In: *2016 IEEE/ACM International Conference on Ad-*

- vances in Social Networks Analysis and Mining (ASONAM)*. IEEE. 2016, pp. 661–668.
- [237] Yannis Panagis and Urška Šadl. “The force of EU case law: A multi-dimensional study of case citations”. In: *Legal Knowledge and Information Systems*. IOS Press, 2015, pp. 71–80.
 - [238] György Görög and Péter Weisz. “Legal entity recognition in an agglutinating language and document connection network for EU Legislation and EU/Hungarian Case Law”. In: *arXiv preprint arXiv:1907.12280* (2019).
 - [239] Dafne van Kuppevelt and Gijs van Dijck. “Answering legal research questions about dutch case law with network analysis and visualization”. In: *Legal knowledge and information systems*. IOS Press, 2017, pp. 95–100.
 - [240] Zihan Huang et al. “Context-aware legal citation recommendation using deep learning”. In: *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law*. 2021, pp. 79–88.
 - [241] Lars Wissler et al. “The Gold Standard in Corpus Annotation”. In: *IEEE GSC, Passau*. June 2014.
 - [242] James Pustejovsky and Amber Stubbs. *Natural Language Annotation for Machine Learning: A guide to corpus-building for applications*. O’Reilly Media, Inc., 2012.
 - [243] Ron Artstein and Massimo Poesio. “Inter-coder agreement for computational linguistics”. In: *Computational Linguistics* 34.4 (2008), pp. 555–596.
 - [244] Donna Spencer. *Card Sorting: Designing Usable Categories*. 1st ed. Rosenfeld Media, Apr. 2009. ISBN: 1933820020.
 - [245] Linton C. Freeman. “A Set of Measures of Centrality Based on Betweenness”. In: *Sociometry* 40.1 (Mar. 1977), pp. 35–41. ISSN: 00380431.

- [246] Manlio De Domenico et al. “Ranking in interconnected multilayer networks reveals versatile nodes.” In: *Nature communications* 6 (2015), p. 6868.
- [247] M. Berlingerio et al. “Foundations of Multidimensional Network Analysis”. In: *2011 International Conference on Advances in Social Networks Analysis and Mining*. 2011, pp. 485–489.
- [248] Sandra G. Hart and Lowell E. Staveland. “Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research”. In: *Human Mental Workload*. Ed. by Peter A. Hancock and Najmedin Meshkati. Vol. 52. Advances in Psychology. North-Holland, 1988, pp. 139–183.
- [249] Steven Bird, Ewan Klein, and Edward Loper. *Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit*. Beijing: O’Reilly, 2009. ISBN: 978-0-596-51649-9. DOI: <http://my.safaribooksonline.com/9780596516499>. URL: <http://www.nltk.org/book>.
- [250] Mathieu Bastian, Sebastien Heymann, and Mathieu Jacomy. “Gephi: An Open Source Software for Exploring and Manipulating Networks”. In: *Third international AAAI conference on weblogs and social media*. The AAAI Press, Mar. 2009, pp. 361–362.
- [251] Wim Peters and Adam Wyner. “Legal text interpretation: Identifying hohfeldian relations from text”. In: *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*. 2016, pp. 379–384.
- [252] C. Biagioli et al. “Automatic Semantics Extraction in Law Documents”. In: *Proceedings of the 10th International Conference on Artificial Intelligence and Law*. ICAIL ’05. Bologna, Italy: Association for Computing Machinery, 2005, pp. 133–140. ISBN: 1595930817.
- [253] Marie-Francine Moens et al. “Automatic Detection of Arguments in Legal Texts”. In: *Proceedings of the 11th International Conference on Artificial Intelligence and Law*. ICAIL ’07. Stanford, California:

- Association for Computing Machinery, 2007, pp. 225–230. ISBN: 9781595936806.
- [254] Harry Surden. “Machine learning and law”. In: *Washington Law Review* 89 (2014), p. 87.
 - [255] Parlamento Europeo e Consiglio dell’Unione Europea. *Regolamento (CE) n. 141/2000 sui medicinali orfani*. Accessed: 2024-10-20. 2000. URL: <https://eur-lex.europa.eu/legal-content/IT/TXT/HTML/?uri=CELEX:32000R0141&from=IT>.
 - [256] Alexander Clark, Chris Fox, and Shalom Lappin. *The handbook of computational linguistics and natural language processing*. John Wiley & Sons, 2013.
 - [257] Karen Sparck Jones. “A statistical interpretation of term specificity and its application in retrieval”. In: *Journal of documentation* (1972).
 - [258] Jari Saramäki et al. “Generalizations of the clustering coefficient to weighted complex networks”. In: *Physical Review E* 75.2 (2007), p. 027105.
 - [259] Ulrik Brandes and Daniel Fleischer. “Centrality measures based on current flow”. In: *Annual symposium on theoretical aspects of computer science*. Springer, 2005, pp. 533–544.
 - [260] Ronald S Burt. “Structural holes and good ideas”. In: *American journal of sociology* 110.2 (2004), pp. 349–399.
 - [261] Gergely Palla et al. “Uncovering the overlapping community structure of complex networks in nature and society”. In: *nature* 435.7043 (2005), pp. 814–818.
 - [262] David Liben-Nowell and Jon Kleinberg. “The Link Prediction Problem for Social Networks”. In: *Proceedings of the Twelfth International Conference on Information and Knowledge Management. CIKM ’03*. New Orleans, LA, USA: Association for Computing Machinery, 2003, pp. 556–559. ISBN: 1581137230.

- [263] Vito Latora and Massimo Marchiori. “Efficient behavior of small-world networks”. In: *Physical review letters* 87.19 (2001), p. 198701.
- [264] Abdol–Hossein Esfahanian. “Connectivity algorithms”. In: *Topics in structural graph theory*. Cambridge University Press, 2013, pp. 268–281.
- [265] Ameet V. Joshi. *Machine Learning and Artificial Intelligence*. Springer, 2020.
- [266] Fabian Pedregosa et al. “Scikit-learn: Machine learning in Python”. In: *the Journal of machine Learning research* 12 (2011), pp. 2825–2830.
- [267] Aric Hagberg, Pieter Swart, and Daniel S Chult. *Exploring network structure, dynamics, and function using NetworkX*. Tech. rep. Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2008.
- [268] Thomas Tullis and Larry Wood. “How Many Users Are Enough for a Card-Sorting Study?” In: *Learning*. June 2004, pp. 1–10.
- [269] Vincent D Blondel et al. “Fast unfolding of communities in large networks”. In: *Journal of statistical mechanics: theory and experiment* 2008.10 (2008), P10008.
- [270] Stefano Civitarese. “A European Convergence towards a Stare Decisis Model”. In: *Rev. Digital de Derecho Admin.* 14 (2015), p. 173.
- [271] Thijmen Koopmans. “Stare decisis in European law”. In: *Essays in European law and integration* (1982), pp. 1957–1982.
- [272] Wolfgang Schön. “Tax legislation and the notion of fiscal aid: A review of 5 years of European Jurisprudence”. In: *State aid law and business taxation* (2016), pp. 3–26.
- [273] Emilio Sulis et al. “Exploiting Textual Similarity Techniques in Harmonization of Laws”. In: *AIxIA 2021 – Advances in Artificial Intelligence*. Ed. by Stefania Bandini et al. Cham: Springer International Publishing, 2022, pp. 185–197. ISBN: 978-3-031-08421-8.

- [274] EUR-Lex. *Access to European Union law*. Accessed: 2024-10-20. 2024. URL: <https://eur-lex.europa.eu>.
- [275] European Union Publication Office. URL: <https://op.europa.eu/en/web/eu-vocabularies/formex/> (visited on 06/09/2022).
- [276] Lionel Neville Brown and Tom Kennedy. *The court of justice of the European communities*. Sweet & Maxwell Uk, 2000.
- [277] Sergey Brin and Lawrence Page. “The anatomy of a large-scale hypertextual web search engine”. In: *Computer networks and ISDN systems* 30.1-7 (1998), pp. 107–117.
- [278] Mark E Dickison, Matteo Magnani, and Luca Rossi. *Multilayer social networks*. Cambridge University Press, 2016.
- [279] European Commission. *Joint Practical Guide of the European Parliament, the Council and the Commission for Persons Involved in the Drafting of European Union Legislation*. 2016.
- [280] Tom Van Engers et al. “Ontologies in the legal domain”. In: *Digital Government: E-Government Research, Case Studies, and Implementation* (2008), pp. 233–261.
- [281] Valentina Leone, Luigi Di Caro, and Serena Villata. “Taking Stock of Legal Ontologies: A Feature-Based Comparative Analysis”. In: *Artificial Intelligence and Law* 28.2 (2020), pp. 207–235. DOI: [10.1007/s10506-019-09252-1](https://doi.org/10.1007/s10506-019-09252-1).
- [282] W3C. *Open Digital Rights Language (ODRL)*. Accessed: 2024-10-20. n.d. URL: <https://www.w3.org/TR/odrl-vocab/>.
- [283] W3C Community Group. *ODRL Community Group*. Accessed: 2024-10-20. n.d. URL: <https://www.w3.org/community/odrl/>.
- [284] Simon Steyskal and Axel Polleres. “Defining expressive access policies for linked data using the ODRL ontology 2.0”. In: *SEM ’14*. 2014.

- [285] Ontology Engineering Group. *Linked Data Rights (LDR)*. Accessed: 2024-10-20. n.d. URL: <http://oeg-dev.dia.fi.upm.es/licensius/static/ldr/>.
- [286] Ontology Engineering Group. *Ontology Engineering Group Homepage*. Accessed: 2024-10-20. n.d. URL: <http://www.oeg-upm.net/>.
- [287] W3C. *Creative Commons Rights Expression Language (ccREL)*. Accessed: 2024-10-20. n.d. URL: <https://www.w3.org/Submission/ccREL/>.
- [288] L4LOD. *Licence for Linked Open Data (L4LOD)*. Accessed: 2024-10-20. n.d. URL: <http://ns.inria.fr/l4lod/v2/l4lod%20v2.html>.
- [289] Distinto et al. *LOTED2 Ontology*. Accessed: 2024-10-20. n.d. URL: <https://code.google.com/archive/p/loted2/source>.
- [290] Isabella Distinto, Mathieu d'Aquin, and Enrico Motta. "LOTED2: An Ontology of European Public Procurement Notices". In: *Semant. Web* 7.3 (Jan. 2016), pp. 267–293. ISSN: 1570-0844. DOI: [10.3233/SW-140151](https://doi.org/10.3233/SW-140151). URL: <https://doi.org/10.3233/SW-140151>.
- [291] TED. *Tenders Electronic Daily (TED)*. Accessed: 2024-10-20. n.d. URL: <https://ted.europa.eu/TED/main/HomePage.do>.
- [292] José Félix Muñoz-Soro et al. "PPROC, an ontology for transparency in public procurement". In: *Semantic Web* 7 (2016), pp. 295–309.
- [293] Muñoz-Soro et al. *Public Procurement Ontology (PPROC)*. Accessed: 2024-10-20. n.d. URL: <http://contsem.unizar.es/def/sector-publico/pproc.html>.
- [294] Harshvardhan Jitendra Pandit et al. "GDPRtEXT - GDPR as a Linked Data Resource". In: *ESWC*. 2018.
- [295] Pandit et al. *GDPRtEXT (GDPR text extensions)*. Accessed: 2024-10-20. n.d. URL: <https://openscience.adaptcentre.ie/ontologies/GDPRtEXT/deliverables/docs/index-en.html>.

- [296] Monica Palmirani et al. “PrOnto: Privacy Ontology for Legal Reasoning”. In: *Electronic Government and the Information Systems Perspective*. Ed. by Andrea Kö and Enrico Francesconi. Cham: Springer International Publishing, 2018, pp. 139–152. ISBN: 978-3-319-98349-3.
- [297] Alessandro Oltramari et al. “PrivOnto: A semantic framework for the analysis of privacy policies”. English (US). In: *Semantic Web 9.2* (2018), pp. 185–203. ISSN: 1570-0844. DOI: [10.3233/SW-170283](https://doi.org/10.3233/SW-170283).
- [298] Usable Privacy Policy Project. *Usable Privacy Policy Project*. Accessed: 2024-10-20. 2018. URL: <https://www.usableprivacy.org/>.
- [299] Publications Office of the European Union. *EuroVoc*. Accessed: 2024-10-20. n.d. URL: <https://publications.europa.eu/en/web/eu-vocabularies/th-dataset/-/resource/dataset/eurovoc>.
- [300] Palmirani et al. and Athan et al. *LegalRuleML*. Accessed: 2024-10-20. n.d. URL: <http://docs.oasis-open.org/legalruleml/legalruleml-core-spec/v1.0/legalruleml-core-spec-v1.0.html>.
- [301] Monica Palmirani et al. “LegalRuleML: XML-Based Rules and Norms”. In: *RuleML America*. 2011.
- [302] Tara Athan et al. “LegalRuleML: Design Principles and Foundations”. In: *Reasoning Web. Web Logic Rules: 11th International Summer School 2015, Berlin, Germany, July 31- August 4, 2015, Tutorial Lectures*. Ed. by Wolfgang Faber and Adrian Paschke. Cham: Springer International Publishing, 2015, pp. 151–188. ISBN: 978-3-319-21768-0. DOI: [10.1007/978-3-319-21768-0_6](https://doi.org/10.1007/978-3-319-21768-0_6). URL: https://doi.org/10.1007/978-3-319-21768-0_6.
- [303] European Legislation Identifier (ELI). *European Legislation Identifier*. Accessed: 2024-10-20. n.d. URL: <https://publications.europa.eu/en/web/eu-vocabularies/th-dataset/-/resource/dataset/eurovoc>.

europa.eu/en/web/eu-vocabularies/model/-/resource/dataset/eli.

- [304] Gandon et al. *Normative Requirements Vocabulary (NRV)*. Accessed: 2024-10-20. n.d. URL: <http://ns.inria.fr/nrv/v1/nrv%20v1.html>.
- [305] Fabien Gandon, Guido Governatori, and Serena Villata. “Normative Requirements as Linked Data”. In: Dec. 2017. DOI: [10.3233/978-1-61499-838-9-1](https://doi.org/10.3233/978-1-61499-838-9-1).
- [306] Alain Auger and Caroline Barrière. “Pattern-based approaches to semantic relation extraction: A state-of-the-art”. In: *Terminology* 14.1 (2008), pp. 1–19.
- [307] Alisa Smirnova and Philippe Cudré-Mauroux. “Relation extraction using distant supervision: A survey”. In: *ACM Computing Surveys (CSUR)* 51.5 (2018), pp. 1–35.
- [308] Amir Pouran Ben Veyseh et al. “Exploiting the syntax-model consistency for neural relation extraction”. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 2020.
- [309] Xikun Zhang et al. “GreaseLM: Graph REASoning Enhanced Language Models”. In: *International Conference on Learning Representations*. 2021.
- [310] A Vaswani. “Attention is all you need”. In: *Advances in Neural Information Processing Systems* (2017).
- [311] Jacob Devlin et al. “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding”. In: *North American Chapter of the Association for Computational Linguistics*. 2019. URL: <https://api.semanticscholar.org/CorpusID:52967399>.
- [312] OpenAI et al. *GPT-4 Technical Report*. 2024. arXiv: [2303.08774](https://arxiv.org/abs/2303.08774) [cs.CL].

- [313] Hugo Touvron et al. “LLaMA: Open and Efficient Foundation Language Models”. In: *ArXiv* abs/2302.13971 (2023). URL: <https://api.semanticscholar.org/CorpusID:257219404>.
- [314] Gemini Team et al. *Gemini: A Family of Highly Capable Multimodal Models*. 2024. arXiv: [2312.11805 \[cs.CL\]](https://arxiv.org/abs/2312.11805). URL: <https://arxiv.org/abs/2312.11805>.
- [315] “The Claude 3 Model Family: Opus, Sonnet, Haiku”. In: URL: <https://api.semanticscholar.org/CorpusID:268232499>.
- [316] Humza Naveed et al. “A Comprehensive Overview of Large Language Models”. In: *ArXiv* abs/2307.06435 (2023). URL: <https://api.semanticscholar.org/CorpusID:259847443>.
- [317] Zhongxiang Sun. “A short survey of viewing large language models in legal aspect”. In: *arXiv preprint arXiv:2303.09136* (2023).
- [318] Jinqi Lai et al. “Large language models in law: A survey”. In: *arXiv preprint arXiv:2312.03718* (2023).
- [319] Lexroom.ai. *Revolutionizing Legal Research: Lexroom.ai’s AI-Driven Solution*. <https://go-dive.net/revolutionizing-legal-research-lexroom-ais-ai-driven-solution/>. Accessed: 28.09.2024. 2024.
- [320] Varun Magesh et al. “Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools”. In: *CoRR* abs/2405.20362 (2024). DOI: [10.48550/ARXIV.2405.20362](https://doi.org/10.48550/ARXIV.2405.20362). arXiv: [2405.20362](https://arxiv.org/abs/2405.20362). URL: <https://doi.org/10.48550/arXiv.2405.20362>.
- [321] Neel Guha et al. “Legalbench: A collaboratively built benchmark for measuring legal reasoning in large language models”. In: *Advances in Neural Information Processing Systems* 36 (2024).
- [322] Ivan L. Padjen. “Systematic Interpretation and the Re-Systematization of Law: The Problem, Co-Requisites, a Solution, Use”. In: *International Journal for the Semiotics of Law - Revue Internationale de Sémiotique Juridique* 33.1 (2020), pp. 189–213. DOI: [10.1007/s11196-019-09672-x](https://doi.org/10.1007/s11196-019-09672-x).

- [323] Rinke Hoekstra et al. “The LKIF Core Ontology of Basic Legal Concepts.” In: *LOAIT* 321 (2007), pp. 43–63.
- [324] Daniela Tiscornia. “The LOIS project: Lexical ontologies for legal information sharing”. In: *Proceedings of the V Legislative XML Workshop*. Citeseer. 2006, pp. 189–204.
- [325] Gianmaria Ajani et al. “The European legal taxonomy syllabus: a multi-lingual, multi-level ontology framework to untangle the web of European legal terminology”. In: *Applied Ontology* 11.4 (2016), pp. 325–375.
- [326] Llio Humphreys. “Populating Legal Ontologies Using Information Extraction Based on Semantic Role Labeling and Text Similarity”. PhD thesis. University of Luxembourg, 2016.
- [327] Luigi Di Caro. “What’s in a Definition? An Investigation of Semantic Features in Lexical Dictionaries.” In: *Proceedings of the 16th International Conference on Web Information Systems and Technologies (WEBIST 2020)*. SCITEPRESS – Science and Technology Publications. 2020, pp. 225–232. ISBN: 978-989-758-478-7. DOI: [10.5220/0010116902250232](https://doi.org/10.5220/0010116902250232).
- [328] Llio Humphreys et al. “Populating legal ontologies using semantic role labeling”. In: *Artificial Intelligence and Law* 29.2 (2021), pp. 171–211.
- [329] Kara Combs et al. “Analogical reasoning: An algorithm comparison for natural language processing”. In: *Proceedings of the 55th Hawaii International Conference on System Sciences*. 2022.
- [330] Víctor Rodríguez-Doncel et al. “A Linked Term Bank of Copyright-Related Terms.” In: *JURIX*. 2015, pp. 91–100.
- [331] Stanford University. *WebProtégé*. Accessed: 2024-10-20. n.d. URL: <https://webprotege.stanford.edu/>.
- [332] Víctor Rodríguez-Doncel. *A Linked Term Bank of Copyright-Related Terms*. Accessed: 2024-10-20. 2015. URL: <http://www.cosasbuenas.es/blog/copyright-term-bank>.

- [333] EasyRDF. *EasyRdf - Converter*. Accessed: 2024-10-20. n.d. URL: <https://www.easyrdf.org/converter>.
- [334] European Data Journalism Network. *Infringeeye. Pending infringement procedures by country and policy area*. <https://www.europeandatajournalism.eu/resources-for-journalists/infringeeye/>. Accessed: 30.09.2024. 2021.
- [335] European Commission. *2023 Annual Report on Monitoring the Application of EU Law*. European Commission, 2023. URL: https://commission.europa.eu/law/application-eu-law/implementing-eu-law/infringement-procedure/2023-annual-report-monitoring-application-eu-law_en.
- [336] Zhikai Chen et al. “Exploring the Potential of Large Language Models (LLMs) in Learning on Graphs”. In: *ACM SIGKDD Explorations Newsletter* 25 (Mar. 2024), pp. 42–61. DOI: [10.1145/3655103.3655110](https://doi.org/10.1145/3655103.3655110).
- [337] Rohan Nanda, Luigi Di Caro, and Guido Boella. “A Text Similarity Approach for Automated Transposition Detection of European Union Directives”. In: *International Conference on Legal Knowledge and Information Systems*. 2016. URL: <https://api.semanticscholar.org/CorpusID:32838837>.
- [338] Rohan Nanda et al. “A unifying similarity measure for automated identification of national implementations of european union directives”. In: *Proceedings of the 16th edition of the International Conference on Artificial Intelligence and Law*. ACM. 2017, pp. 149–158.
- [339] Tadas Klimas and Jurate Vaiciukaite. “The Law of Recitals in European Community Legislation”. In: *ILSA Journal of International and Comparative Law* 15 (2008), pp. 61–93. URL: <https://api.semanticscholar.org/CorpusID:143144202>.
- [340] Abhimanyu Dubey et al. “The llama 3 herd of models”. In: *arXiv preprint arXiv:2407.21783* (2024).

- [341] Ziqi Jin and Wei Lu. “Tab-CoT: Zero-shot Tabular Chain of Thought”. In: *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*. Ed. by Anna Rogers, Jordan L. Boyd-Graber, and Naoaki Okazaki. Association for Computational Linguistics, 2023, pp. 10259–10277. DOI: [10.18653/v1/2023.FINDINGS-ACL.651](https://doi.org/10.18653/v1/2023.FINDINGS-ACL.651). URL: <https://doi.org/10.18653/v1/2023.findings-acl.651>.
- [342] Jason Wei et al. “Chain of Thought Prompting Elicits Reasoning in Large Language Models”. In: *ArXiv* abs/2201.11903 (2022).
- [343] *Neo4J website*. <https://neo4j.com/>. Accessed: 2024-09-30.
- [344] Galileo Sartor et al. “Automated extraction and representation of citation network: a CJEU case-study”. In: *International Conference on Conceptual Modeling*. Springer. 2022, pp. 102–111.
- [345] Michael Hicks, James Humphries, and Joe Slater. “ChatGPT is bullshit”. In: *Ethics and Information Technology* 26 (June 2024), pp. 1–10. DOI: [10.1007/s10676-024-09775-5](https://doi.org/10.1007/s10676-024-09775-5).