Detecting Contractual Risk in ODRL Policies Using DPV

Elena Molino-Peña[1], Jos De Roo[2], Beatriz Esteves[2], José María García[1], and Antonio Ruiz-Cortés[1]

[1] SCORE Lab, I3US Institute, Universidad de Sevilla, Sevilla, Spain
{mmolino,josemgarcia,aruiz}@us.es
[2] IDLab, Department Electronics and Information Systems, Ghent University – imec, Ghent, Belgium
{jos.deroo,beatriz.esteves}@ugent.be

Abstract

Terms of Service (ToS) in Software as a Service (SaaS) ecosystems define the responsibilities of providers and customers. However, the monitoring and enforcement of these contracts is usually under the control of providers, creating risks of power asymmetry and limited transparency. Moreover, many clauses cannot be easily verified, operationalised or computationally interpreted in an unequivocal manner, making it difficult for customers to determine whether an obligation has been fulfilled, a right can be effectively exercised, or a prohibition has been violated, leading to a lack of awareness of the risks associated with accepting such terms. In this paper, a semantic approach is presented to model and infer contractual risks derived from potentially unfair terms in ToS. The Open Digital Rights Language (ODRL) is used to represent contractual terms, while the Data Privacy Vocabulary (DPV) is used to describe the risks to which customers may be exposed when subscribing to SaaS services. We introduce a ToS risk specification to classify risk sources and mitigation measures associated with unfair terms. In addition, we describe a set of N3 rules to infer risks from ODRL policies and link them to DPV risk instances. The approach is implemented in a prototype and illustrated using abusive and non-abusive test cases.

Keywords: ODRL, Terms of Service, Contractual Risk Detection, DPV

1 Introduction

Digital services are typically governed by Terms of Service (ToS), adhesion contracts drafted unilaterally by providers in natural language that define the conditions under which services are offered, including the rights and obligations of the parties, renewal provisions, changes, indemnities, and termination terms, among other legal aspects [11]. These legal documents are usually published as HTML or PDF files, limiting interoperability and preventing automated reasoning [16].

As a result, the verifiability of their terms remains uncertain. Objectively assessing whether parties comply with the agreement is challenging, particularly when customers cannot independently determine whether the provider is fulfilling its commitments [2]. In such situations, contractual compliance is often determined by the provider, effectively granting them discretionary authority, or in the event of a dispute, by the body responsible for resolving the case. The lack of mechanisms increases the risk of consumer vulnerability, reduces transparency, and increases power imbalance. When such clauses are also potentially unfair, meaning that they create a significant imbalance in the parties' rights and obligations to the detriment of the consumer [1], they become a source of contractual risk [4].

A recent evaluation by the European Commission of the EU consumer protection framework, the Fitness Check on Digital Fairness [4], has identified the widespread presence of potentially

unfair terms in ToS, particularly in online transactions. The assessment highlights that the current regulatory framework is not fully adapted to the digital environment. In online contexts, consumers tend to pay less attention, rarely read contractual terms, and often fail to process all the information provided. They are also more vulnerable to digital persuasion techniques, such as dark patterns and other potentially unfair practices, which may expose them to additional risks. This situation produces a significant economic impact, with the annual detriment suffered by consumers estimated at approximately €7.9 billion, whereas the estimated regulatory compliance cost for businesses is around €0.6 billion [4].

To address this problem, this paper explores a semantic approach to formally represent and infer contractual risks associated with potentially unfair terms. Building on legal foundations such as the Unfair Contract Terms Directive (UCTD) and the Fitness Check on Digital Fairness, we investigate the use of N3 inference rules to derive risk instances described using the Data Privacy Vocabulary (DPV) from contractual policies represented in the Open Digital Rights Language (ODRL). The contributions of this study are as follows.

- A ToS Risk specification that introduces a vocabulary to classify risks associated with unfair terms in contractual clauses. It includes 1) risk source classes for unfair terms, and 2) mitigation measures for legal control.
- A set of N3 inference rules to derive risk from ODRL policies. The rules detect potentially unfair terms in the policy and infer corresponding DPV risk instances with mitigation measures.
- A prototype implementation illustrating the approach through abusive and non-abusive clause test cases.

The remainder of the article is organised as follows. Section 2 presents the legal background, relevant policy and data protection specifications, and related work. Section 3 introduces our approach to detecting risks in ToS. Section 4 describes the proposed ToS specification and inference approach. Section 5 presents the validation of the proposal, and finally Section 6 concludes the article.

2 Background and Related Work

2.1 Legal Basis for Unfair Terms

EU consumer law operates under the premise that "the consumer is in a weaker position vis-à-vis the seller or supplier, as regards both his bargaining power and his level of knowledge"[1], and therefore requires protection through public regulation. In this section, we present key legal concepts for consumer protection in digital services derived from the Unfair Contract Terms Directive (UCTD) [1] and the Digital Fairness Fitness Check of EU consumer law [4], focusing on risks and unfairness in digital contracts. The EU Digital Fairness Act[2], currently under development, aims to improve consumer protection and may introduce new terminology and requirements.

The **Unfair Contract Terms Directive (UCTD)** provides a legal framework to protect consumers against unfair terms in contracts concluded between a supplier and a consumer. According to Art. 3, a term that has not been individually negotiated may be considered

[1] https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:62010CJ0618
[2] https://commission.europa.eu/law/law-topic/consumer-protection-law/review-eu-consumer-law_en

unfair if it causes a significant imbalance in the parties' rights and obligations to the detriment of the consumer. The Annex of the Directive provides a non-exhaustive list of contractual terms that may be regarded as potentially unfair, including provisions that limit the supplier's liability or the consumer's legal remedies, such as compensation, in cases of non-performance or harm (Annex (a), (b)), exclude consumers' right to take legal action by requiring arbitration (Annex (q)), allow the supplier to alter or terminate the contract or service unilaterally without valid reason or prior notice (Annex (j), (k), (g)), or bind the consumer to terms they had no real opportunity to review before concluding the contract (Annex (i)). Clauses not listed in the UCTD Annex may also be considered unfair. For instance, clauses limiting the supplier's liability for damage or removal of consumer data may fall within the scope of the Annex (b). Similarly, clauses assigning jurisdiction to foreign courts or determining the applicable law may hinder consumers' ability to exercise their legal rights and may therefore fall under Annex (q) [12, 3].

The **Digital Fairness Fitness Check of EU consumer law** is a European Commission assessment of consumer protection law in the digital environment. In Annex VI, Section VI.1.6 (Unfair contract terms) highlights that the lack of transparency and fairness in ToS may materialise risks in the form of financial harm, lost time, privacy harmed, unintentionally provide personal data, losing access to the service and difficulties with exercising consumer rights. These risks can be related to each type of unfair contractual term. Arbitration, Choice of Law and Jurisdiction may create the potential risk of excluding or hindering the consumer's right to take legal action. Limitation of Liability may place the entire risk on the consumer, without compensation, leading to loss of data or service, reputation and economic loss. Unilateral Change creates legal uncertainty, lack of notice, exclusion of legal rights and harm. Unilateral Termination may result in loss of access to the service or the account and consumer detriment. Contract by using can cause financial harm due to hidden or inaccessible terms. Finally, Content removal may involve loss or damage of consumer data, prohibition to recover stored data, and omission of information.

2.2 Policies and Risk Modelling

ODRL [9] is a W3C standard model for expressing policies, offering a flexible interoperable modelling language and common vocabulary [8] to represent statements about the use of digital content and services. The core model of ODRL includes essential concepts such as permissions, prohibitions, and duties, as well as actions, assets, parties, and constraints, which are necessary to define the responsibilities typically found in ToS. ODRL supports extensibility through its profiling mechanism, enabling the incorporation of additional terms for specific domains. In this context, the TOSL profile [14] incorporates specific elements of online contracts, such as jurisdiction, applicable law, and liability, thereby enabling automated analysis tasks such as the detection of potentially unfair terms [16] and facilitating the representation of complex relationships present in these contracts.

From the perspective of risk modelling, DPV [18] is a W3C specification that provides a set of common concepts for representing and exchanging information about the processing of personal and non-personal data, as well as related technologies such as cloud services and AI. In addition to its core vocabulary, DPV includes several extensions, including the risk extension[3], which introduces terms for describing threats, vulnerabilities, impacts, consequences, and mitigations associated with data processing activities.

The detection of potentially unfair terms in ToS has been widely studied in the context

[3] https://w3id.org/dpv/risk

of consumer protection. Prior research has mainly focused on the automated detection of unfair contract terms using machine learning techniques, and explaining why such clauses are considered potentially unfair under existing legal frameworks [11, 13, 5, 10]. Complementary work has addressed the detection of dark patterns in online services and the risks associated with manipulative practices [17, 7, 19]. These studies highlight how digital service environments can influence user behaviour and potentially lead to consumer harm.

3 Risk Detection in Terms of Service

Figure 1 illustrates the overall workflow of the proposed approach. Contractual clauses are formalised as ODRL policy agreements, including permissions, obligations and prohibitions, which are represented as a knowledge graph (KG). The KG is analysed using N3 reasoning rules to perform rule-based inference, identifying KG nodes corresponding to potentially unfair or non-compliant policy terms, and inferring the associated DPV-structured risk instances. These risk instances capture the source of the risk by linking it to the corresponding ODRL rule, as well as the potential consequences, impacts on consumers, and recommended legal mitigation measures. The result is an enriched KG that integrates both the original ODRL policy representation and the inferred DPV risk instances.

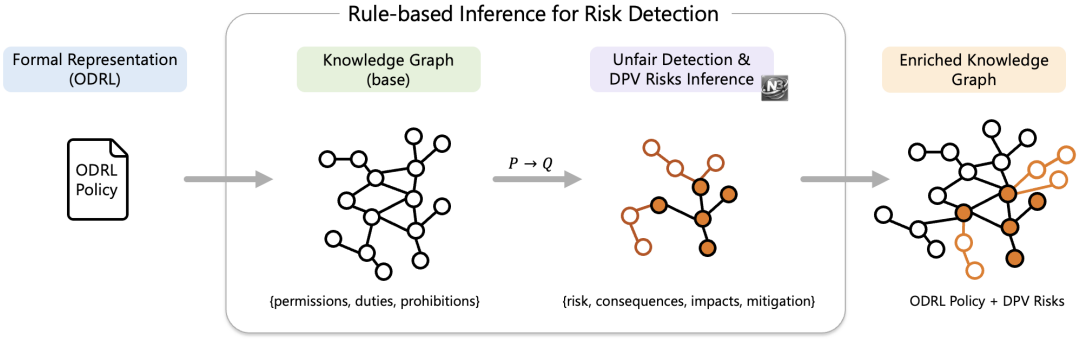

Figure 1: Overview of the proposed approach pipeline

As an example, consider the following clause from Uber's terms of use[4]: "15.3. Uber may, in its sole discretion, terminate these Terms or discontinue the App(s)/Website(s) Services at any time by giving you advance notice in writing." This clause may be considered potentially unfair under Annex (g) of the UCTD, as it does not specify the serious grounds that would justify termination, nor does it define what constitutes a reasonable notice period in terms of the number of days. The clause can be modelled in ODRL following the structure outlined in Listing 1. Granting a termination right over the agreement asset without limitations may introduce several risks for the consumer, as it could result in loss of access to data, service interruptions, or even additional harms, such as financial losses.

In DPV, a risk is defined as the possibility or potential for negative events to occur. The immediate result of a risk materialising is the consequence, which generally affects the system or service. When such consequences result in harm to the consumer, they are classified as impacts. Mitigation measures can be specified to address these risks. However, the measures

[4] https://www.uber.com/legal/en/document/?name=general-terms-of-use&country=belgium&lang=en

considered in our model do not directly mitigate the technical consequences or impacts on the consumer, such as financial loss or data loss. Instead, they mitigate the contractual and legal risks associated with potentially unfair terms, improving transparency and increasing provider responsibilities. Consumers may adopt additional preventive measures to reduce the risks to which they are exposed, such as data backups or provider backups.

Listing 1: TOSL–DPV–ODRL representation of an Uber Unilateral Termination clause

```
:permissionTerminateTerms a odrl:Permission ;
   dcterms:description "Provider can terminate the agreement and services at their discretion";
   odrl:assignee :provider ;
   odrl:action tosl:terminate ;
   odrl:target :service ;
   odrl:duty [
       a odrl:Duty ;
       odrl:assignee :provider ;
       odrl:action odrl:inform ;
       odrl:informedParty :customer ;
       odrl:constraint [
           a odrl:Constraint ;
           odrl:leftOperand odrl:deliveryChannel ;
           odrl:operator odrl:isA ;
           odrl:rightOperand :customerContact
           ]
       ] .

:provider a dpv-odrl:ServiceProvider .
:customer a dpv-odrl:Customer .
:service a tosl:Service .
:customerContact a pd:Contact .
```

Based on Table 1, in the previous example of Uber's terms of use clause, customers assume risks such as data unavailability and service termination, which could be legally mitigated by incorporating constraints and duties into the Terms, such as termination notices with a reasonable advance period (e.g., defined in days or months) and strong justification requirements. Listing 2 illustrates how these risks can be modelled using DPV.

Listing 2: DPV Risks representation of the Unilateral Termination clause

```
:R1UT a dpv:Risk ;
   skos:broader risk:DataUnavailable, risk:ServiceTermination ;
   risk:hasRiskSource [
       a tosl-risk:UnilateralTerminationClause ;
       dpv:hasContractualClause :permissionTerminateTerms
   ] ;
   dpv:isMitigatedByMeasure tosl-risk:TerminationNotice,
                                tosl-risk:ProvideJustification ;
   dpv:hasConsequence [
      a risk:DataLoss, risk:ServiceNotProvided ;
      dpv:hasImpact [
          a risk:FinancialLoss, risk:NonMaterialDamage, risk:CustomerConfidenceLoss, risk:Harm ;
          dpv:hasImpactOn :customer
      ]
   ] .

:permissionTerminateTerms a dpv:Permission .
:customer a dpv-odrl:Customer .
```

4 A DPV-based Framework for Modelling Risks of Unfair Terms

This section outlines a risk modelling specification to represent risks arising from potentially unfair contractual terms contained in ODRL policies, using concepts from DPV. The main goal is to support the modelling of such risks in order to enable automatic risk assessment and the identification of legal mitigation strategies, thereby improving the semantic representation of the ToS and reducing risks for consumers. The specification is informed by the legal provisions discussed in Section 2. Unfair clauses are modelled as `risk:RiskSource`, representing sources of risk associated with specific contractual terms, while legal mitigation mechanisms are represented as `dpv:RiskMitigationMeasure`. The defined specification is available at https://w3id.org/tosl/risk.

4.1 Terms of Service Language Risk Specification

The semantic representation of risk is provided in DPV through its risk management extension[5], which follows the ISO 31000 standard. The vocabulary includes several concepts to support risk assessment and management, including `dpv:Risk`, `dpv:RiskMitigationMeasure`, `dpv:Consequence`, `dpv:Impact`, `dpv:hasLikelihood`, `dpv:hasSeverity`, and `dpv:hasRiskLevel`. In addition, it provides a taxonomy including concepts such as `risk:FinancialLoss`, `risk:Harm`, and `risk:DataBreach`.

These concepts enable the structured representation of risks as illustrated in Listing 2. For each risk source, the model captures the potential associated risks, their possible immediate consequences and impacts on consumers, as well as some legal mitigation measures aimed at reducing unfairness. Thus, we can model the potential harms that may arise from the rights and obligations defined in the ToS. However, DPV does not provide terms to associate ODRL rules with the specific type of unfair term that constitutes the risk source. Furthermore, it does not include legal mitigation concepts that can be associated with suggested ODRL rules and constraints to enrich the analysed ODRL policy.

Building on the DPV specifications and the legal findings described in Section 2, we provide a set of terms to represent unfair terms as `risk:RiskSource` and legal mitigation mechanisms as `dpv:RiskMitigationMeasure`. The specification introduces a parent concept `tosl-risk:UnfairContractClause`, which serves as an abstraction for clauses that may generate risks for consumers. From this concept, we define subclasses corresponding to each of the specific types of unfair clause sources derived from the UCTD [1] and previous works [11, 13, 4], including `tosl-risk:ArbitrationClause`, `tosl-risk:UnilateralChangeClause`, `tosl-risk:UnilateralTerminationClause` and `tosl-risk:LimitationOfLiabilityClause`, among others.

Regarding the definition of legal mitigation measures, the UCTD provides the legal basis for determining when contractual terms may be considered potentially unfair. These provisions serve as the foundation for defining a set of mitigation terms that can be linked to ODRL rule representations. For instance, limitation of liability clauses may be less abusive if consumers retain access to legal remedies, such as compensation for damages caused by service disruptions. This mitigation measure is represented through the concept `tosl-risk:RightToCompensation`, which aims to capture the consumer's right to receive compensation when damages occur. Similarly, unilateral change and termination clauses may be mitigated by providing prior notice within a defined time period (`tosl-risk:TerminationNotice`) and by specifying grounded

[5] https://w3id.org/dpv/risk

Table 1: Terms for defining DPV risk instances based on potentially unfair terms.

Risk Source[a]	Risks	Consequences	Impacts	Mitigations
A	RightsDenied RightsLimited RightsObstructed	LegalSupportLimited RightsExercisePrevented InabilityToEstablishLegal-Claims	JudicialCosts	AccessToCourts ConsumerResidence-Jurisdiction ConsumerCountryGoverning-Law
LAW	RightsDenied RightsLimited RightsObstructed	LegalSupportLimited RightsExercisePrevented	JudicialCosts	ConsumerCountryGoverning-Law
J	RightsDenied RightsLimited RightsObstructed	LegalSupportLimited RightsExercisePrevented	JudicialCosts	ConsumerResidence-Jurisdiction
CR	DataUnavailable DataErasureError	DataLoss	NonMaterialDamage FinancialLoss CustomerConfidenceLoss	DeletionNotice DeletionJustification
USE	RightsLimited RightsObstructed	QualityDegraded	NonMaterialDamage FinancialLoss CustomerConfidenceLoss	ExplicitConsent-Requirement
LTD	ServiceLimited DataUnavailable Harm	DataLoss ReputationalLoss	FinancialLoss RightsDenied RightsLimited	RightToCompensation
CH	ServiceCostIncreased ServicePartiallyProvided	QualityDegraded ServiceQualityReduced ServiceProvisionDelayed ServiceNotProvided	NonMaterialDamage FinancialLoss CustomerConfidenceLoss Harm	ChangeNotice ChangeJustification ReaffirmConsent
TER	DataUnavailable ServiceTermination	DataLoss ServiceNotProvided	NonMaterialDamage FinancialLoss CustomerConfidenceLoss Harm	TerminationNotice TerminationJustification

[a] This nomenclature was defined by Lippi et al. and by Micklitz et al. [11, 13]. A stands for Arbitration, LAW for Choice of Law, J for Jurisdiction, CR for Content Removal, USE for Contract by Using, LTD for Limitation of Liability, CH for Unilateral Change, and TER for Unilateral Termination.

reasons for the action (`tosl-risk:TerminationJustification`). Furthermore, arbitration clauses may also be mitigated by ensuring that consumers retain the right to take legal action without being required to submit disputes to arbitration, represented through the concept `tosl-risk:AccessToCourts`.

The remaining mitigation terms are listed in Table 1, which also outlines the relationship between potentially unfair terms and their associated risk instances using DPV concepts. Risks, consequences and impacts are existing concepts in the DPV specification, which were mapped to unfair terms in previous work [15] based on legal sources such as the Fitness Check Report. In contrast, the risk source and mitigation taxonomies are defined in our specification, described in more detail at https://w3id.org/tosl/risk.

4.2 N3 rules for Risk Detection

Notation3 (N3) [20] is an RDF-based language that provides constructs for expressing rules and performing reasoning over knowledge graphs. N3 rules are expressed as logical implications between graph patterns and can be evaluated by N3 reasoners such as Eyeling [6]. In our approach, rules are used to detect contractual patterns in ODRL policies that may correspond to potentially unfair terms, and to infer the associated contractual risks. A set of rules has

been developed, one for each type of unfair term, and is available on GitHub[6].

The rules are applied to ODRL policies represented as RDF graphs, such as the permission excerpt shown in Listing 1. A policy graph is a finite set of RDF triples describing the deontic contractual terms, including permissions, duties and prohibitions together with their associated constraints. Reasoning is performed through N3 implication rules of the form $P \rightarrow Q$, as shown in Listing 3. When the antecedent (P) of a rule matches a pattern in the policy graph, the rule is triggered and the triples defined in the consequent (Q) are inferred. In our setting, the antecedent is responsible for identifying ODRL rules that match patterns considered potentially unfair, such as ODRL rules that grant capabilities without the necessary constraints or associated duties. The consequent generates the corresponding DPV risk instances. This allows each inferred risk to be traced back to the policy rule that triggered it.

Listing 3: Conceptual pattern for risk detection in ODRL Policies

```
{
    ?G log:includes { ?R } .
    ?G log:notIncludes { ?S1 ?S2 ?SN } .
} => {
    ?risk a dpv:Risk .
} .
```

More precisely, the antecedent of the rule is composed of (1) a positive inclusion condition that filters policy rules (`log:includes`), where `?R` describes the rule pattern, such as a permission for the provider to change service features or remove user content; and (2) a finite set of negative inclusion conditions that detect the absence of specific safeguards (`log:notIncludes`), where `?S1`, `?S2`, and `?SN` represent safeguard patterns, such as reaffirmed consent, justification requirements, or prior notification. For instance, Listing 4 presents the antecedent of the Unilateral Termination N3 rule. This rule filters the permissions where the provider is the assignee, the action is `terminate`, and the asset is the service, while excluding cases that include justification and prior notification constraints.

Listing 4: Antecedent of the N3 rule for Unilateral Termination clauses

```
{
    ?agreement :policyGraph ?G .
    ?G log:includes {
        ?perm a odrl:Permission ;
        odrl:assignee ?p ;
        odrl:action tosl:terminate ;
        odrl:target ?t .

        ?p a tosl:Provider .
        ?t a tosl:Service .
    } .
    ?G log:notIncludes {
        ?perm odrl:constraint [
            odrl:leftOperand tosl:justification
        ] .
        ?perm odrl:duty [
            odrl:action odrl:inform
        ] .
        ?perm odrl:constraint [
            odrl:leftOperand odrl:delayPeriod
        ] .
    } .
}
```

[6] https://github.com/isa-group/risk-e-tos/tree/main/rules

Listing 5: Consequent of the N3 rule for Unilateral Termination clauses

```
{
    ?risk a dpv:Risk ;
        skos:broader risk:DataUnavailable, risk:ServiceTermination ;
        risk:hasRiskSource [
            a tosl-risk:UnilateralTerminationClause ;
            dpv:hasContractualClause ?perm
        ] ;
        dpv:isMitigatedByMeasure tosl-risk:TerminationNotice,
                                    tosl-risk:TerminationJustification,
                                    tosl-risk:RightToCompensation ;
        dpv:hasConsequence [
           a risk:DataLoss, risk:ServiceNotProvided ;
           dpv:hasImpact [
               a risk:FinancialLoss, risk:NonMaterialDamage,
                   risk:CustomerConfidenceLoss, risk:Harm ;
               dpv:hasImpactOn :customer
           ]
        ] .

    tosl-risk:TerminationNotice
        :suggestAdd {
            ?perm
            odrl:duty [
                a odrl:Duty ;
                odrl:assignee :provider ;
                odrl:action odrl:inform ;
                odrl:informedParty :customer ;
                odrl:constraint [
                    a odrl:Constraint ;
                    odrl:leftOperand odrl:deliveryChannel ;
                    odrl:operator odrl:isA ;
                    odrl:rightOperand :customerContact
                ];
            ] .
        } .

    tosl-risk:TerminationJustification
        :suggestAdd {
            ?perm
            odrl:constraint [
                a odrl:Constraint ;
                odrl:leftOperand dpv-odrl:Justification ;
                odrl:operator odrl:isAnyOf ;
                odrl:rightOperand (
                       justifications:SecurityImpaired
                       justifications:UnlawfulActivityObjection
                       justifications:LegalObligation
                )
            ] .
        } .

    tosl-risk:RightToCompensation
    :suggestAdd {
        ?perm
        odrl:duty [
            a odrl:Duty ;
            odrl:assignee :provider ;
            odrl:action odrl:compensate ;
            odrl:informedParty :customer ;
            odrl:refinement [
                a odrl:Constraint ;
                odrl:leftOperand odrl:payAmount;
                odrl:operator odrl:eq;
                odrl:rightOperand "3000"^^xsd:decimal;
                odrl:unit <http://dbpedia.org/resource/Euro>
            ]
        ] .
    } .
}
```

The consequent consists of a set of pre-defined DPV risks and ODRL-based mitigation suggestions, based on Table 1. These risks are instantiated and associated with the identified unfair policy rule, as illustrated in Listing 5. The inferred DPV triples identify the detected unfair term as the source of the risk (`risk:hasRiskSource`), specify its potential consequences (`dpv:hasConsequence`) and associated impacts (`dpv:hasImpact`), and include a set of recommended mitigation measures (`dpv:isMitigatedByMeasure`). These mitigations are expressed as ODRL constraints and duties that can be incorporated into the original policy rule to partially mitigate the identified risk in line with the UCTD, as discussed in Section 2.

For instance, Listing 5 presents the consequent of the Unilateral Termination N3 rule, where risks such as `risk:DataUnavailable` and `risk:ServiceTermination` are linked to the detected termination permission. Each mitigation measure is accompanied by an ODRL representation that proposes the suggested modifications. For instance, `TerminationNotice` introduces a duty requiring the provider to inform the customer through an appropriate communication channel. `TerminationJustification` adds a constraint specifying acceptable justifications for termination, while `RightToCompensation` introduces a duty requiring the provider to compensate the customer.

5 Validation

To demonstrate the feasibility of the proposed approach, a prototype has been developed that allows the execution of N3 rules over ODRL policies using the Eyeling reasoner [6]. The prototype is available online[7] together with the set of rules[8] used for risk detection. To evaluate the rules, a test suite of contractual clauses was created. The dataset includes examples of both abusive and non-abusive clauses for each type of potentially unfair term. These policies are expressed as ODRL agreements, and represent the common ToS clauses such as unilateral termination, limitation of liability and contract by using. The system analyses the policy graph and identifies patterns corresponding to potentially unfair terms. When such patterns are detected, the associated risks are inferred. As output, a structured representation of the inferred risk is produced using the DPV model and returned in Turtle format. Each risk instance includes the associated risk source, along with its consequences and impacts, and a set of recommended mitigation measures modelled using ODRL. These mitigation measures are intended to refine the corresponding ODRL rule in order to reduce the unfairness of the clause. The coverage of the system depends on the set of implemented rules and the availability of machine-readable policies. Additional work is required to expand the rule set and to evaluate the approach on larger collections of real-world ToS.

6 Conclusions

Potentially unfair terms included in ToS can create significant risks for consumers. This paper introduces the ToS Risk specification and defines N3 inference rules to derive such risks from ODRL policies as DPV risk instances. These contributions are demonstrated through a prototype capable of analysing both abusive and non-abusive clauses, illustrating how ODRL policy representations can support the identification of risks and mitigation measures. As future work, we plan to extend the current rules and move towards a more user-centered risk analysis. This

[7] https://github.com/isa-group/risk-e-tos
[8] https://github.com/isa-group/risk-e-tos/tree/main/rules

would allow risks to be assessed according to specific user requirements, where users could provide relevant constraints, such as their country of residence, whether they permit the deletion or sharing of their data, or whether they require prior notice before changes to the agreement, among others. The aim is the identification of risks that can be ranked and explained based on user constraints, as well as comparisons across providers.

Acknowledgments

This work was supported by the R&D projects PID2021-126227NB-C21, PID2021-126227NB-C22, and PDC2022-133521-I00, funded by MICIU/AEI/10.13039/501100011033/ERDF/EU; TED2021-131023B-C21, funded by MICIU/AEI/10.13039/501100011033/European Union NextGenerationEU/PRTR; and PID2024-155693NB-C44, PDC2025-166101-C21 and PDC2025-166101-C22 funded by MICIU/AEI/10.13039/501100011033/FEDER, EU; TSI-100930-2023-2, funded by Ministerio de Asuntos Económicos y Transformación Digital; and DGP_PIDI_2024_00918 funded by Consejería de Universidad, Investigación e Innovación, Junta de Andalucía.

References

[1] COUNCIL DIRECTIVE 93/13/EEC of 5 April 1993 on unfair terms in consumer contracts. *Official Journal of the European Communities*, 1993.

[2] Fernando Castillo, Eduardo Brito, Sebastian Werner, Pille Pullonen-Raudvere, and Jonathan Heiss. Towards Trusted Service Monitoring: Verifiable Service Level Agreements. In *Service-Oriented Computing – 23rd International Conference, ICSOC 2025, Shenzhen, China, December 1–4, 2025, Proceedings, Part II*, 2026.

[3] European Commission. Guidance on the interpretation and application of Council Directive 93/13/EEC on unfair terms in consumer contracts, 2019.

[4] European Commission. Fitness Check of EU Consumer Law on Digital Fairness. Commission Staff Working Document SWD(2024) 230 final, European Commission, Brussels, October 2024.

[5] Sławomir Dadas, Marek Kozłowski, Rafał Poświata, Michał Perełkiewicz, Marcin Białas, and Małgorzata Grebowiec. A support system for the detection of abusive clauses in B2C contracts. *Artificial Intelligence and Law*, 2024.

[6] Jos De Roo. Eyeling: A compact Notation3 (N3) reasoner in JavaScript, 2026. v1.22.16. https://zenodo.org/records/20784774.

[7] Colin M. Gray, Cristiana Santos, and Nataliia Bielova. Towards a Preliminary Ontology of Dark Patterns Knowledge. In *CHI Conference on Human Factors in Computing Systems*, 2023.

[8] R. Iannella, M. Steidl, S. Myles, and V. Rodríguez-Doncel. ODRL Vocabulary & Expression 2.2. W3C Recommendation, W3C, 2018. [Accessed 13-03-2026].

[9] R. Iannella and S. Villata. ODRL Information Model 2.2. W3C Recommendation, W3C, 2018. [Accessed 13-03-2026].

[10] S. Kokila and D. Kerena Hanirex. Effective ML Approach for Identification of Unlawful Practices in Online Terms of Service. In *Intelligent Cyber Physical Systems and Internet of Things (ICoICI)*, 2024.

[11] Marco Lippi, Przemysław Pałka, Giuseppe Contissa, Francesca Lagioia, Hans-Wolfgang Micklitz, Giovanni Sartor, and Paolo Torroni. CLAUDETTE: an automated detector of potentially unfair clauses in online terms of service. *Artificial Intelligence and Law*, 27(2):117–139, 2019.

[12] Marco Loos and Joasia Luzak. Update the Unfair Contract Terms Directive for Digital Services. Technical Report PE 676.006, Policy Department for Citizens' Rights and Constitutional Affairs, Directorate-General for Internal Policies, February 2021.

[13] Hans-W. Micklitz, Przemysław Pałka, and Yannis Panagis. The Empire Strikes Back: Digital Control of Unfair Terms of Online Services. *Journal of Consumer Policy*, 40(3):367–388, 2017.

[14] Elena Molino-Peña, José María Cruz-Lorite, José María García, and Antonio Ruiz-Cortés. TOSL: An Ontology to Detect Abusive Services. In *Joint Proceedings of the ESWC 2025 Workshops and Tutorials co-located with 22nd Extended Semantic Web Conference (ESWC 2025), Portorož, Slovenia, June 1–2, 2025*, 2025.

[15] Elena Molino-Peña, Beatriz Esteves, José María García, and Antonio Ruiz-Cortés. Towards a Risk-Based Analysis of Terms of Service for Consumer Decision Support. In *4th Privacy & Personal Data Management Session @ Solid Symposium*, 2026.

[16] Elena Molino-Peña, José María García, and Antonio Ruiz-Cortés. Integrating Terms of Service and Service Level Agreements for Automating Cloud Service Management. In *Service-Oriented Computing – 23rd International Conference, ICSOC 2025, Shenzhen, China, December 1–4, 2025, Proceedings, Part II*, 2026.

[17] Apichaya Nimkoompai. Risk Analysis of Encountering Dark Patterns of UX E-commerce Applications Affecting Personal Data. In *6th International Conference on Information Technology*, 2022.

[18] Harshvardhan J. Pandit, Beatriz Esteves, Georg P. Krog, Paul Ryan, Delaram Golpayegani, and Julian Flake. Data Privacy Vocabulary (DPV) – Version 2.0. In *The Semantic Web – ISWC 2024*, 2025.

[19] Cristiana Santos, Viktorija Morozovaite, and Silvia De Conca. No harm no foul: how harms caused by dark patterns are conceptualised and tackled under EU data protection, consumer and competition laws. *Information & Communications Technology Law*, 34(3):329–375, 2025.

[20] W. Van Woensel, D. Arndt, P.-A. Champin, D. Tomaszuk, and G. Kellogg. Notation3 Language. Technical report, W3C, 2023. [Accessed 13-03-2026].