

A Corpus for Multilingual Analysis of Online Terms of Service

Kasper Drazewski¹  and Andrea Galassi²  and Agnieszka Jablonowska³ 

Francesca Lagioia³  and Marco Lippi⁴  and Hans Wolfgang Micklitz³ 

Giovanni Sartor³  and Giacomo Tagiuri³ and Paolo Torroni² 

¹Bureau Européen des Unions de Consommateurs, Brussels, Belgium

²DISI, University of Bologna, Italy

³European University Institute, Florence, Italy

⁴DISMI, University of Modena and Reggio Emilia, Italy

^{1,3}{name.surname}@eui.eu

²a.galassi@unibo.it, paolo.torroni@unibo.it

⁴marco.lippi@unimore.it

Abstract

We present the first annotated corpus for multilingual analysis of potentially unfair clauses in online Terms of Service. The data set comprises a total of 100 contracts, obtained from 25 documents annotated in four different languages: English, German, Italian, and Polish. For each contract, potentially unfair clauses for the consumer are annotated, for nine different unfairness categories. We show how a simple yet efficient annotation projection technique based on sentence embeddings could be used to automatically transfer annotations across languages.

1 Introduction

Cultural and linguistic diversity is a guiding principle of the European Union. Within the context of an ever expanding EU, the concept of multilingualism stands out as one of the most prominent symbols of European historical, political, social and cultural diversity.¹ From a legal perspective, the EU commitment to multilingualism is significant not only as a guarantee of legal certainty, clarity, and transparency, but also democratic accountability. Accordingly, for example, the EU legislation is generally published in all official languages.²

While multilingualism certainly has important benefits, it also comes with costs and challenges.

¹A recent indication of the importance of multilingualism for the EU is found in the European Commission Communication of 22 November 2005 “A New Framework Strategy for Multilingualism” (COM(2005) 596 final).

²Regulation No 1 determining the languages to be used by the European Economic Community [1958] OJ 17/385 with subsequent amendments. Consider also the right to good administration, as expressed in Article 41(4) of the Charter of Fundamental Rights of the European Union [2012] OJ C 326/391.

To illustrate, the existence of 24 official versions of every legal act, all of which are equally authentic,³ may inevitably create interpretative difficulties. In this regard, the Court of Justice has repeatedly found that the wording used in one language version of a Union provision cannot serve as the only basis for its reading, and has pointed to the role of cross-language comparisons as well as teleological and systematic methods of interpretation.⁴ Ethnic languages, and related requirements, are also an important issue in the EU internal market and consumer protection policies. The EU has so far mostly left this matter to national authorities, recognizing the associated tensions and cultural contexts.⁵ For example, the EU has refrained, so far, from regulating languages in which standard terms in consumer contracts should be drafted, allowing for differing approaches to emerge in various jurisdictions (Loos, 2017).⁶ In view of the above, and considering that consumer protection authorities and non-governmental organisations in Europe tend to operate in their respective languages, ef-

³Judgment of the Court of 6 October 1982, C-283/81 - CILFIT, ECLI:EU:C:1982:335, para. 18.

⁴See, e.g., judgment of the Court of 19 December 2013, C-281/12 - Trento Sviluppo, ECLI:EU:C:2013:859, para. 25 and the case law cited.

⁵On the one hand, language requirements give rise to additional costs for the traders who wish to engage in cross-border transactions. On the other hand, they might be justified for consumer protection reasons.

⁶In the context of pre-contractual information see also: Article 6(7) of Directive 2011/83/EU of the European Parliament and of the Council of 25 October 2011 on consumer rights, amending Council Directive 93/13/EEC and Directive 1999/44/EC of the European Parliament and of the Council and repealing Council Directive 85/577/EEC and Directive 97/7/EC of the European Parliament and of the Council [2011] OJ L 304/64.

fective consumer protection technologies must be capable of dealing with multilingual landscapes.

In this work, we focus on the automated detection of unfair clauses in Terms of Service (ToS), which are usually available in multiple languages. In particular, we investigate methods for automatically transferring the annotations made on ToS in the context of the Claudette project⁷ onto the corresponding versions of the same documents in a target language, where such resources and expertise may be lacking.

Our ultimate goal is to use automatically generated annotations for training linguistic tools, usually available only in English, like Claudette, for the target language without resorting to expert annotators in that language. This would leverage the creation of classifiers that can leverage the linguistic resources available in the target language, to analyze documents in that language.

From a machine learning point of view, the question is whether we shall necessarily train independent models for each and every language. This paper only partially answers the question, by extending previous work (Galassi et al., 2020) towards three different languages and by making a novel multilingual corpus available to the community. Future research will consider training and comparing the classifiers.

2 Corpus

We built a multilingual parallel corpus consisting of 25 Terms of Service annotated in English, Italian, German and Polish. These languages are spoken in large EU countries as well as in different regions, and they have been selected based on the availability of mother tongue legal experts for the annotation task.

The analysed ToS were retrieved from the Claudette pre-existing corpus, covering 100 English ToS (Lippi et al., 2019; Ruggeri et al., 2021). Such terms mainly concern popular digital services provided to consumers, including leading online platforms (such as search engines and social media). The predominant language of drafting of these ToS is English, with differing availability of corresponding ToS in other languages.

To carry out the present study, the ultimate 25 ToS⁸ were selected on the basis of three main cri-

teria: a) their availability in the four selected languages; b) the possibility of identifying a correspondence between the different versions, given their publication date; and c) the similarity of their structure (e.g. number of clauses, sections, etc.). To illustrate, while ToS in both German and Italian were identified for 63 out of the 100 ToS contained in the pre-existing Claudette training corpus, Polish versions were found for only 42 of these 63 ToS. Out of the 42 ToS available in the four languages, we selected those with the more closely corresponding versions based on criteria b) and c) above. Perfect correspondence across the 4 languages, however, could not be achieved for all 25 ToS. As further discussed in Section 3, some discrepancies may persist.

We hypothesise that a particularly strong similarity between the English ToS and other language versions exists where the latter are translations of the former. Although we could not assess this comprehensively in the present study, we infer from the wording of the ToS that at least in 9 out of 25 cases, German, Italian and Polish documents were indeed translations of the English originals. More specifically, in the ToS of 9 companies⁹ we identified explicit references to the language of the terms. Consider, for example, the following clause taken from Mozilla ToS (updated June 23, 2016) and its corresponding German version:

In the event of a conflict between a translated version of these Terms and the English language version, the English language version shall control. (line 71).

Im Falle eines Konflikts zwischen der übersetzten Version dieser Bedingungen und der englischen Fassung gilt die englische Fassung. (line 71).

A more comprehensive verification whether different language versions of the ToS were translated from the English original was not possible at this stage of research.

The analysed corpus contains, in its English version, 6,831 sentences and 678 labelled clauses, divided in nine categories of clause unfairness, i.e., arbitration (<a>), unilateral change (<ch>), content removal (<cr>), jurisdiction (<j>), choice of law

Ryanair, Skype, Skyscanner, Snapchat, Spotify, Terravision, Tinder, Tripadvisor, Tumblr, Uber, Weebly, Yelp and Zynga.

⁹Booking.com, Dropbox, Evernote, Grindr, Mozilla, Skype, Tumblr, Yelp, Zynga.

⁷<http://claudette.eui.eu>

⁸In particular, we selected the ToS offered by: Booking.com, Dropbox, Electronic Arts, Evernote, Facebook, Garmin, Google, Grindr, LinkedIn, Mozilla, Pinterest, Quora,

(<law>), limitation of liability (<ltd>), unilateral termination (<ter>), contract by using (<use>), and privacy included (<pinc>). While the first eight categories have been described by Lippi et al. (2019), the latter, i.e., the privacy included category, has been recently introduced to identify clauses stating that consumers consent to the privacy policy simply by using the service. Such clauses have been always considered potentially unfair.¹⁰

In the annotations, to indicate the degree of unfairness a numeric value has been appended to each XML tag, with a value 1 meaning clearly fair, 2 potentially unfair, and 3 clearly unfair. For instance, the <j3> label indicates that a certain clause concerns jurisdiction and it has been marked as clearly unfair.

Table 1 shows some statistics on the corpus, by reporting the number of annotated clauses for each tag, across the four different languages. The whole annotated corpus is made freely available for research purposes and can be downloaded at http://claudette.eui.eu/corpus_multilingual_NLLP2021.zip.

3 Discrepancies

Our analysis revealed some discrepancies across the different language versions of the selected documents. This makes the projection task particularly challenging, also affecting the performance evaluation whenever a perfect match between documents is missing.

A general observation concerns the asymmetric length of documents across the different language versions. In some cases, this is the product of missing/extra clauses, for example because of country-specific terms. This is the case, among others, for the Terravision and Spotify ToS (see discussion infra in this section). In other cases, asymmetric length is a feature of linguistic differences. Take for example Evernote ToS that, while showing a remarkable similarity of structure across the 4 languages (no extra/missing clauses), shows a different word and sentence count across the 4 languages: 6,798 words for the Polish version (283 sentences), 7,580 for the German (291 sentences),

¹⁰As an example consider the following clause taken from the Evernote ToS: <pinc2>By using the Service, you acknowledge, accept and agree with all provisions of the Privacy Policy, including, without limitation, the use and treatment of the text, images, and other data you choose to input, upload, or store in Evernote (collectively, “Content”) and your personal information in accordance with such Privacy Policy.</pinc2>.

	DE	EN	IT	PL
a1	3	3	4	4
a2	22	29	29	35
a3	3	4	4	4
ch2	98	100	103	103
ch3	1	1	0	0
cr2	25	27	28	26
cr3	23	24	24	26
j1	14	15	15	15
j3	46	49	48	50
law1	18	16	16	19
law2	33	39	36	36
ltd1	19	27	16	17
ltd2	212	229	216	229
ltd3	1	1	1	1
pinc2	17	21	20	21
ter2	69	71	71	75
ter3	49	49	50	49
use2	54	58	58	61
total	707	753	739	771

Table 1: Corpus statistics: we report the number of annotated clauses for each tag, across the four different languages. Suffices 1, 2, and 3 represent levels of fairness: 1 means clearly fair, 2 stays for potentially unfair, and finally 3 for clearly unfair.

7,902 for the Italian (270 sentences) and 7,795 for the English version (279 sentences).

Moving from the structure of documents to the analysis of sentences, we identified four types of discrepancy across languages.

The first one concerns sentence structure. This type of discrepancy is mostly illustrated by situations in which the same information, contained in a single clause in the English version, is split into more than one clause in the target languages (or vice versa). In some cases such situations do not appear very problematic as the same tag of the original English sentence is identically reproduced across the two sentences in the target language.¹¹

In other cases the implications for annotation are more problematic, as nested tags may need to be manually split. Consider for instance the following sentence from the Tripadvisor ToS (updated October 1, 2018) in its English and Italian version.

<ch2><use2>We may change or otherwise modify the Agreement in the future

¹¹See, e.g., the clause on contract termination in English and Italian versions of the Dropbox ToS (updated April 17, 2018), line 58.

in accordance with the Terms and Conditions herein, and you understand and agree that your continued access or use of this Website after such change signifies your acceptance of the updated or modified Agreement. (line 7).

TripAdvisor si riserva il diritto di cambiare o altrimenti modificare l'Accordo in futuro, in conformità ai termini e alle condizioni ivi espresse. *L'Utente accetta pertanto che l'accesso a o l'utilizzo continuato del portale successivamente alle modifiche implica l'accettazione dell'Accordo così come aggiornato o modificato.* (line 7).

Other problematic cases may be those in which only one of the two sentences in the target language is relevant and should be marked as unfair. Consider the following clause from the Google ToS (updated April 30, 2014) and their Polish version:

We may review content to determine whether it is illegal or violates our policies, and we may remove or refuse to display content that we reasonably believe violates our policies or the law. (line 21)

Możemy sprawdzać treści w celu określenia, czy są one nielegalne lub czy naruszają nasze zasady. Jeśli stwierdzimy, że dany materiał jest niezgodny z tymi zasadami lub obowiązującym prawem, możemy go usunąć lub odmówić jego wyświetlenia. (line 21)

In this case, the different segmentation results in only one of the sentences in the target language being annotated as a potentially unfair content removal clause. The preceding part, related to the reviewing of the content, remains irrelevant for the envisaged task.

A second type of discrepancy consists in errors or inaccuracies in translation from the source English ToS into the target languages. Consider the following example from the English and Polish versions of the Yelp ToS (updated November 27, 2012).

We reserve the right to modify, update, or discontinue the Site at our sole discretion, at any time, for any or no reason, and without notice or liability. (line 141)

Zastrzegamy sobie prawo do modyfikowania, aktualizowania i wyłączenia Witryny wedle własnego uznania, w dowolnym momencie, z dowolnego powodu i bez wcześniejszego powiadomienia. (line 141)

While the first clause was correctly marked under three categories, the Polish version only relates to termination and unilateral change, without any reference to liability. Such a discrepancy does not exist in either German or Italian, which suggests a translation inaccuracy.

Similarly, consider the following clause taken from the Evernote ToS and their Italian version (updated August 28, 2017):

If you do register for or otherwise use our Service you shall be deemed to confirm your acceptance of the Terms and your agreement to be a party to this binding contract. (line 4).

La registrazione al Servizio o il suo uso richiedono di confermare la propria accettazione dei Termini, cosa che rende di conseguenza l'utente un soggetto vincolato al presente contratto. (line 4).

Here the English sentence is a clear example of a contract by using clause, thus potentially unfair (by registering, the consumer consents to the ToS). By contrast, the Italian translation simply states that in order to register for the service, the consumer needs to accept the terms. Thus, the Italian version has not been annotated. A similar discrepancy is not observed in the German and Polish versions of this sentence, which again suggests a translation error.

The third type of discrepancy we identify is missing/extra clauses, meaning the absence of certain clauses in the different versions of the ToS. Consider, for instance, the following clause taken from the Terravision ToS (date of update not provided; ToS accessed on 31 January 2021), which is completely missing in the Italian ToS:

<lt2>Terravision is not liable for any lost or stolen items or baggage.</lt2>
(line 101).

The same is true for arbitration and jurisdiction clauses in the English version of the Spotify ToS, which are absent in the German version (effective as of 9 September 2015).¹²

Terms of service in target languages may further include extra (potentially) unfair clauses. Consider, for example, the Polish version of the Garmin ToS (updated April 3, 2014) containing an extra clause marked under the category of privacy included.¹³

Missing/extra clauses may be the product of an omission or error, but they also may be product of deliberate choices to draft contracts differently for different national markets. Such different choices may be made necessary by country-specific services offered only at certain locations, by the need to comply with country-specific regulation, or, more generally, as product of the company's intention to regulate the contractual relationship differently in different countries. For example, the need to comply with country specific regulation may result in substantive differences in clauses providing for age requirements to use certain services.¹⁴ In a few cases, especially in the German versions of some ToS, an entirely different approach is adopted to the drafting of clauses related to certain matters, like provider's liability. Among others such discrepancies were found in the Electronic Arts ToS.¹⁵

The fourth type of discrepancy derives from the choice of legal terminology, which may be more or less deliberate. Indeed, legal languages remain profoundly culture-bound and terminological nuances in different languages may be difficult to capture regardless of the applied translation process.

In our corpus this type of discrepancy emerges in relation to the use of particular terms, for instance terms related to responsibility and liability. Unlike in English, in the other languages of our corpus such a distinction is hard to capture.

Consider, for instance, the following clause taken from the Pinterest ToS (updated May 1, 2018) and

¹²More specifically, a sizeable proportion of ToS concerning dispute resolution was removed from the German version and replaced with a single clause (line 175).

¹³Cf. the Polish version of Garmin ToS, line 74.

¹⁴Cf. the English and German versions of Electronic Arts ToS (updated May 17, 2018) at line 33.

¹⁵See the English version of Electronic Arts ToS at line 164, and the German version of corresponding ToS at lines 164-166.

its Polish and Italian counterparts that employ the expressions “odpowiedzialność” and “responsabilità”, without differentiation between responsibility and liability:

<lt2>Pinterest takes no responsibility and assumes no liability for any User Content that you or any other person or third party posts or sends using our Service.</lt2> (line 88)

<lt2>Pinterest nie ponosi odpowiedzialności za jakiegokolwiek Treści użytkownika, które Ty lub inni użytkownicy albo osoby trzecie umieszczacie lub wysyłacie przy użyciu naszych usług.</lt2> (line 88)

<lt2>Pinterest non si assume alcuna responsabilità per i Contenuti dell'utente che tu o qualsiasi altra persona o terza parte pubblicare o inviare utilizzando il nostro Servizio.</lt2> (line 88)

While linguistic variations related to responsibility (obligations) and liability do not always entail differing annotations, they sometimes may. Consider, for example, the following situation from the Terravision ToS (as accessed on 31 January 2021):

<lt2>Terravision is not liable to passengers who did not reserve their trip.</lt2>

Terravision non ha alcun obbligo nei confronti dei passeggeri che non hanno prenotato la corsa.

Here, a translation of the English term "to be liable" into the more generic Italian "avere obbligo" results in different annotation choices in the different languages.

Finally, as a general remark, it is important to note that deviations from the English source ToS are uneven across languages, being largest in German ToS. This may suggest that the drafting of such ToS is done by human agents, who may pay more attention to the national legal context and specific terminologies than automated translators. For example, note the markedly different take on the matter of privacy and data protection in these clauses of Spotify, where the German drafters refrained from packaging data protection consent with the agreement to the ToS:

<pinc2>Your agreement with us includes these Terms and Conditions of Use (“Terms”) and our Privacy Policy.</pinc2> <pinc2>(The Terms, Privacy Policy, and any additional terms that you agree to, as discussed in the Entire Agreement section, are referred to together as the “Agreements”).</pinc2>
(line 37)

Ihre Vereinbarung mit uns schließt diese Geschäfts- und Nutzungsbedingungen („Bedingungen“) ein sowie jegliche weitere Vereinbarung, der Sie zustimmen, wie im Abschnitt Vollständiger Vertrag beschrieben (gemeinsam als die „Vereinbarungen“ bezeichnet). (line 37)

The hypothesis of a more careful drafting being applied in the drafting of German language documents seems confirmed by the fact that retrieving identical corresponding versions of ToS was most hard for German. Conversely, we observed a lower mismatch in both Polish and Italian ToS, where significant structural differences can be retrieved only in limited cases.

4 Experiments

We employ the same methodology proposed in Galassi et al. (2020), by extending the study to two additional languages, namely Italian and Polish. The general idea of the methodology is to transfer the annotations of a given English document onto the same document, given in a different language. As a first step, we exploit an automatic translation system to obtain English sentences from the target language (i.e., German, Polish, or Italian). Then, we find correspondences between the original (tagged) English sentences and the ones (untagged) obtained via translation. In this way, the annotations of the first document can be associated to the sentences in the second document. For this step, any distance or metric that measures the (dis-)similarity between sentences can be exploited. In this work, we use the distance between ELMo (Peters et al., 2018) sentence embeddings coupled with Dynamic Time Warping (Berndt and Clifford, 1994), which was the approach achieving the best performance in Galassi et al. (2020). Finally, the projected annotations are then easily associated back to the original sentences in the target language.

	DE	IT	PL
Precision	0.87	0.94	0.90
Recall	0.95	0.98	0.97
F1-macro	0.91	0.94	0.91
F1-micro	0.91	0.96	0.93
F1-weighted	0.91	0.96	0.93

Table 2: Projection results for the three languages.

To segment our documents into sentences we exploit the `Spacy` python package,¹⁶ whereas for translation into English we use the `deep_translator` python library, in particular the Google translator.¹⁷ We measure the performance of our system by considering the result of the projection as a multi-label classification task on the sentences of the target documents. We report precision as the percentage of projected tags that are correct and recall as the percentage of original tags that have been correctly matched to the target document, in both cases macro-averaged over the documents. F1-macro is the harmonic mean between the first two metrics. We also report the F1-micro, which sums up all the true positives, false positives, and false negatives, without distinguishing among the classes, i.e., tags. Finally, the F1-weighted is the weighted sum of the F1 scores computed per each class (tag), where the weight is given by the support (i.e., number of tagged sentences) of each tag.

As shown in Table 2, the projection works very well for all three languages, with an F1-macro larger than 0.90 for all the languages. The approach performs particularly well for Italian, where a 0.97 recall is obtained, with a 0.94 precision. Polish is the second-best performing language, with German only slightly worse in terms of precision, which is still very high, though.

5 Discussion

Discrepancies across language versions may significantly affect the performance of the projection task. In particular, we found that the system produces some erroneous results where a) there is no correspondence between clauses (e.g., some of them are missing in one or more language versions); b) sentences are split differently; c) there is a mismatch in the ToS structure; d) incorrect translations or

¹⁶<https://spacy.io/>

¹⁷<https://deep-translator.readthedocs.io/>

specific linguistic choices can be identified.

As an example of no correspondence between clauses, consider the Electronic Arts ToS (updated May 17, 2018), in which liability clauses are built in an entirely different manner between English and German. In this case the projection missed the two following clauses, which have no equivalent in the English version:

<td2>Bei kostenpflichtigen EA-Services ist die Haftung von EA außer in Fällen der Verletzung von Leib, Leben oder Gesundheit auf Vorsatz und grobe Fahrlässigkeit beschränkt.</td2> [...] <td2>In Fällen leichter Fahrlässigkeit haftet EA jedoch nur für die schuldhaftige Verletzung von wesentlichen Vertragspflichten, wie z.B. Verzug oder Unmöglichkeit der Leistung Die Haftung im Falle der Verletzung einer solchen wesentlichen Vertragspflicht ist beschränkt auf den typischen Schaden, der bei Abschluss dieses Vertrages für EA aufgrund der Umstände vorhersehbar war.</td2> (line 162).

As for the different segmentation of sentences, consider the following example taken from the Evernote ToS (updated August 28, 2017), where two consecutive `<ch2>` clauses in the English version were split by a third innocuous clause in the German version.

<ch2>These Terms may be amended as new features, technology, or legal requirements arise, so please check back from time to time.</ch2><ch2>If we make a significant change, we'll notify you and, where required, seek your consent.</ch2>

<ch2>Diese Bedingungen können sich ändern, wenn neue Funktionen, Technologien oder rechtliche Anforderungen vorliegen.</ch2> [misidentified as <ch2>] Werfen Sie also immer mal wieder einen Blick auf diese Seite, um stets auf dem neuesten Stand zu sein. <ch2>Bei wesentlichen Änderung werden wir Sie darüber informieren und gegebenenfalls Ihre Einwilligung einholen.</ch2>

Division of longer passages into shorter sentences as well as discrepancies in the documents' structure have been identified as recurrent causes of false positives and false negatives also in relation to examples reported in Section 3. This is the case, among others, for the previously described term on content removal in the Polish version of the Google ToS, where only one of the resulting sentences in the target language merited annotation. Erroneous projection outcomes could also be identified when a longer English sentence, marked under two categories using a nested tag, was divided into two shorter sentences in the target language, each annotated under one category. The previously mentioned clause in the Tripadvisor ToS on unilateral change and contract by using illustrates such a case. Moreover, a comparably high number of false positives and negatives was observed in relation to Terravision, where the document structure varied in different language versions, among others due to country-specific terms.

Finally, as an example of incorrect projection due to translation errors and linguistic choices, consider the case of a clause in Grindr ToS (updated 13 March 2020), where the phrase "bez zobowiązania" (without commitments) was used instead of a reference to liability. While this clause was not marked in the target language, the system nonetheless incorrectly recognized it as a potentially unfair limitation of liability.

The same was true for the previously mentioned clause in the Italian Terravision ToS, referring to a generic term "obbligo". Erroneous projection results were further identified in relation to the previously mentioned clause in the Yelp ToS, which in the Polish version lacked reference to liability. Similarly, a clause on contract by using, which was marked in the English ToS of Evernote, but not in their Italian version, was found among false positives.

6 Conclusions

This paper presents an extension of the annotation projection technique presented in Galassi et al. (2020) for the automated detection of potentially unfair clauses in online ToS by applying it to multiple languages. The presented approach directly addresses the call for automatically transferring the annotations made on legal documents in a language with significant linguistic resources and domain experts, such as English, onto the corresponding ver-

sions of the same documents in a target language, where such resources and expertise may be lacking. This approach does not apply to online ToS only, but to a wide category of legal documents, and in particular within the European Union, with a great variety of different languages.

For these reasons, we presented the first multilingual corpus for the detection of potentially unfair clauses for the consumers in online ToS. The data set consists of 25 contracts tagged in four different languages (English, German, Italian, and Polish) for a total of 100 documents. We illustrated the discrepancies across the different versions of the same contract, and we analyzed the major causes of errors in the projection process.

In the future, we will further study correspondence between the English and other language versions, and consider whether computational methods can be deployed to confirm existence of translations. Our final aim is to compare machine learning predictors trained on original and projected annotations, to assess whether the projection procedure can substitute the time-consuming task of document annotation performed by experts.

References

- Donald J Berndt and James Clifford. 1994. [Using dynamic time warping to find patterns in time series](#). In *KDD workshop*, pages 359–370. AAAI Press.
- Andrea Galassi, Kasper Drazewski, Marco Lippi, and Paolo Torrioni. 2020. [Cross-lingual annotation projection in legal texts](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 915–926.
- Marco Lippi, Przemysław Pałka, Giuseppe Contissa, Francesca Lagioia, Hans-Wolfgang Micklitz, Giovanni Sartor, and Paolo Torrioni. 2019. [Claudette: an automated detector of potentially unfair clauses in online terms of service](#). *Artificial Intelligence and Law*, 27(2):117–139.
- Marco B.M. Loos. 2017. [Double dutch-on the role of the transparency requirement with regard to the language in which standard contract terms for b2c-contracts must be drafted](#). *Journal of European Consumer and Market Law*, 6(2).
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *NAACL-HLT*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Federico Ruggeri, Francesca Lagioia, Marco Lippi, and Paolo Torrioni. 2021. [Detecting and explaining](#)

[unfairness in consumer contracts through memory networks](#). *Artificial Intelligence and Law*, pages 1–34.