

TOWARDS AUTOMATING INCONSISTENCY CHECKING OF LEGAL TEXTS

Tomer Libal / Tereza Novotná

Assistant Professor, American University of Paris, Computer Science Math and Environmental Science
5 Boulevard de la Tour-Maubourg, 75007 Paris, FR
tlibal@aup.edu; <https://tomer.libal.info>

PhD candidate, Masaryk University, Faculty of Law, Institute of Law and Technology
Veveří 158/70, 611 80 Brno, CZ
tereza.novotna@mail.muni.cz; <https://www.law.muni.cz>

Keywords: *Legal inconsistency, Legal drafting, Legal reasoning, Automated reasoning*

Abstract: *Inconsistencies in legal texts are sometimes hard to find and correct. In this paper a novel machine executable methodology is described, which can facilitate this process. We demonstrate the methodology on the American University of Paris Library Circulation Policy No. AA036EN.*

1. Introduction and related work

As any practicing lawyer would agree, legal texts are as full of indeterminacy, vagueness and ambiguity as any other kind of text. Therefore, law evolved whole sets of interpretative rules to bridge these logical phenomenons. Legal inconsistency is not an exemption, although this kind of legal indeterminacy can, as the authors will show in this article, be sometimes solved using automatic legal logical reasoning. For the purposes of this article, legal indeterminacy is used as a general term for any form of legal indeterminacy containing legal inconsistency as well as legal vagueness and ambiguity.

Legal text inconsistency definition emerges from the logical definition of inconsistency, which is logical contradiction. Logical contradiction is commonly defined as the conjunction of two statements that cannot be both simultaneously true in any possible world. When applied to the legal order, the definition lies in two rules, where first rule demands a certain behaviour, whereas the second rule demands the exact opposite behaviour [ENGEL, 2004]. As such, the inconsistencies can be in many cases solved using the common legal interpretative rules, if we are aware of a hierarchy of the opposite legal rules (*lex specialis*) or their chronological sequence (*lex posterior* rule), although the application of these rules is limited to certain cases with certain conditions (for example the hierarchy of rules).

Legal inconsistencies can appear on the one hand in statements of different stakeholders in legal cases (the inconsistency of facts), but on the other hand they are also very common in normative statutory and regulatory legal texts (the inconsistency of legal rules). As such, they arise from different aims of different stakeholders in legislative procedure and following political and social discussion as well as simply out of the legislature's inattention. This means, among other things, that legal indeterminacy (and inconsistency) is not only a grammatical or syntactic problem, but rather a theoretical-legal question aimed at the very essence of law and its social function.

The discussion concerning legal indeterminacy goes back as far as the 1970s, although it is certainly not its very beginning. In the 1970s and 1980s, legal scholars formed critical legal studies to oppose Dworkin's „No

Right Answer?“ thesis [DWORKIN, 1977]. Critical legal scholars argued that there is always a counterrule for every legal rule and that there is always a certain level of uncertainty in a legal case [TUSHNET, 2005]. This leads to the very question of the legitimacy of the law, which is indeterminate and inconsistent by definition [KRESS, 1989, p. 286]. Although the legal indeterminacy discussion is much less active since the 1990s, lawyers, judges and legal scholars encounter legal inconsistency on a daily basis. However, with the development of artificial intelligence and information technology, the 21st century offers new ways to identify legal indeterminacies and new tools that could fill in this gap and help legal practitioners in working with contradictory legal texts.

Concerning text contradiction, research in this field is mainly focused on detecting contradictions and entailment in the text. However, most state-of-the-art work is focused on the contradiction detection in news and journals. A non-exhaustive list of related work is mentioned below. One of the first attempts in this domain is based on first order logic [COOPER, 1996; CONDORAVDI et al., 2003]. The PASCAL Recognizing Textual Entailment is one of the first attempts to detect textual entailment [DAGAN et al., 2009; BAR-HAIM et al., 2006; GIAMPICCOLO et al., 2007]. Another work which focuses on the detection of contradiction is the one of [HARABAGIU et al., 2006]. This work is based solely on machine learning algorithms and focuses on a special kind of contradiction. From the logical linguistic models for inconsistency detection, [TATU et al., 2007], the BLUE system developed by [CLARK et al., 2009], and a hybrid NatLogsystem by [MACCARTNEY et al., 2009] should be mentioned. All of the above mentioned models are corpora-based or dataset-based meaning that they contain corpora or they build on a machine learning generated one.

In this paper we take another approach. The interpretation of a legal text is translated into logic formulae. Automated theorem provers are then used in order to deduce that the formulae are inconsistent. Moreover, we define a partially automated methodology which can be used to detect inconsistencies in the text and modify it such that the inconsistencies are eliminated. The tool defined in [LIBAL/STEEN, 2019] is well suited for our purpose. It contains an annotation based editor which facilitates the translation of the legal interpretation into logical formulae. In addition, it supports executing automated theorem provers by clicking different buttons in the editor. It doesn't really detect the inconsistency in legal text itself, it has no intention to actually understand the meaning of the content of a legal document. The meaning is only dependent on the user's interpretation of the legal text. Therefore the user is responsible for the correctness of the interpretation for which a behaviour-driven methodology exists [LIBAL/STEEN, 2019]. This allows the user to work with any possible legal text and in any possible language, since the NAI tool only registers the variables the user defines for the natural text. The broad scope of possible uses of the NAI tool is one of the main advantages. Secondly, one of its main goals is to abstract over most of the technical complexities. Although the NAI tool is still in an early development stage, using annotations of legal text, we believe, is an easier task for possible users than direct translations into formal languages. It is necessary to state that further developing the user interface is part of our future work plans. In addition, the NAI tool is fully available through an internet browser and it is open access and free.

In the next Section we formally describe the problem we are trying to solve and we introduce an example which will be used throughout the paper. Section 3. is devoted to the introduction of the NAI tool [LIBAL/STEEN, 2019], which is the technology we use in order to find and correct inconsistencies in legal texts. Section 4. will describe how the tool can actually be used towards this aim using a partially automated methodology. Lastly, we conclude in Section 5. and describe some possible future work.

2. Inconsistency of legal texts

In this paper, we will use the American University of Paris Library Circulation Policy No. AA036EN as an example for detecting inconsistencies in legal documents.¹ The Policy Statement reads as follows:

1. *In order to borrow library items (books, films, CDs, etc.) the user must present a current and valid identification card issued by the University (or an affiliated program).*
2. *Access to the AUP Library and borrowing privileges are for personal use only.*
3. *The AUP ID card is personal and may not be lent to others.*
4. *The ID card issued to AUP students, faculty, staff and alumni requesting an alumni card acts as ID card, library user card and a building access card.*
5. *Patrons from affiliated programs must also carry a special access card (issued by your program office) in addition to their institution's ID card.*
6. *The University Library is not open to the public. For external users, please refer to AA038EN – Library External User.*
7. *Misuse of library privileges can lead to disciplinary action.*

The scenario we use in order to demonstrate the automatic inconsistency checking is the following: a person with a valid and current identification card issued by an affiliated program lends this card to a friend who uses the card for borrowing a library book. The person then gets a disciplinary notice for lending her card. The library notice contains a very brief explanation that the person violated the obligation of personal use. This obligation, according to the library representatives, arises from provisions 2. and 3. of the Policy Statement, which identify personal use as the only possible use of the library services. The person however doesn't agree with the library representatives' reasoning over the Policy Statement. According to her, the conditions for borrowing the library books are mentioned in the first statement. The person acknowledges the additional obligation arising from the fifth sentence – the obligation to carry a special access card, which both she and her friend carry. The second and third statement, on the other hand, indicate the obligation of personal use only for the AUP issued card cardholder. Although the second sentence indicates the general obligation that borrowing books is personal, the third sentence, as *lex specialis*, refines the interpretation to be applicable only on the card issued by AUP, which shall be used only personally. It is obvious there are different ways how to interpret the rules in the Policy and in some of the ways the inconsistency can be solved (e.g. using teleological interpretation, where the aim is to borrow always personally notwithstanding the type of card). However we choose the interpretation that cannot be solved using the interpretative rules in order to demonstrate the application of the NAI tool.

As both interpretations of the subject text are reasonable, one can deduce that the text is inconsistent. It contains two contradictory interpretations according to which the person with a current and valid card from an affiliated party (and special access card) is both permitted and prohibited to lend the card to a friend to use the library services with. It is obvious that the inconsistency is mainly caused by the second and third sentences and the way they are interpreted regarding the rest of the Policy Statement. Using the interpretative methods, one can argue for different legal relationships between sentences 1., 2. and 3. The argument of the library representatives is based on the fact that sentences 1. and 2. are both general conditions and personal use is just another general condition for using its services. Sentence 3. then refines the obligation by

¹ The text of this policy is available at:
<https://www.aup.edu/university-policies-guidelines/library-circulation>

indicating that AUP issued card is personal, but it doesn't itself exclude the card issued by an affiliated party from the obligation in sentence 2. On the other hand, the argument of the reprimanded person suggests that sentence 3. is *lex specialis* in relation to sentences 1. and 2. and thus the prohibition to lend the card to another person concerns only the AUP issued card as it indicates in sentence 3. Both interpretations of the Policy Statement are correct, but at the same time, there is no legitimate situation in which it is possible to at the same time permit and prohibit a behaviour of a person, we can thus state an inconsistency in this legal text.

3. The NAI Tool

For the purpose of finding inconsistencies in legal texts, we use the NAI tool.² The NAI tool is an open source web application for automatic legal reasoning over any legal text. The NAI tool and its functioning is properly described in [LIBAL/STEEN, 2019]. We offer here only a basic overview of the main functions. Using the NAI tool, a user can transform a legal text in natural language into logical formulae. The transformation into logical formulae is called legal formalization. The formalization of legal texts is done through annotations of important legal terms and relations among them in the *Annotation* tab of the NAI tool, as can be seen after logging in to the demo account and selecting a legislation. This formalization task is solely dependent on a user's interpretation of the legal text. The tool then works only with the terms and relations that a user defines during formalization and not with the original text. The user needs to define important parts of the text (or assumptions) as „*terms*“. Terms are variables or first-order terms which denote legally relevant parts of the text. They are interconnected via the use of „*connectives*“. Connectives are logical relations between terms (for example conjunction, disjunction, implication and equivalence as well as permission, obligation and prohibition).

In addition to the *Annotation* tab, the tool contains three other tabs - *Vocabulary*, *Formalization* and *Advanced formalization*. All the annotated terms are stored in the *Vocabulary* tab with their meaning - the text that was used for annotation.

The NAI tool is able to reason automatically over legal texts. That eventually means that it can answer queries whether the facts and solution of a unique case logically follows from a formalization. The NAI tool is also able to check for the consistency of terms and logical relations in the text. This feature is of special interest for the task described in this paper. Both of these tasks can be performed using two types of buttons - „*Execute query*“ and „*Run consistency check*“. The tool then answers that the goal statement either logically follows the assumptions and logical relations of the case and the legal document or the opposite, that it doesn't logically follow the assumptions. In some cases, the tool can timeout and state that it is not able to answer the query. Further versions of NAI will improve over this drawback. Due to its use of automated theorem provers, every answer is supported by a mathematical certificate of correctness. Since these certificates can be further checked by independent programs, the chances of an error are very small.

3.1. Legal formalization of example case

² All the formalization described in this paper is performed in the following special account on the website www.nai.uni.lu, account: iris@nai.lu, password: nai. We hereby encourage reader to see what we are describing directly in this account for better understanding.

³Concerning the example case we introduced in Section 2., the NAI tool is able to detect the inconsistency in the text if the correct formalization of the text is given. It turns out that a relatively simple formalization which uses propositional logic only, suffices for capturing and correcting the inconsistency. The NAI tool is able to work in the same manner with first-order logic as well depending on the complexity of the legal text and the user's priorities. We assume that the following interpretation of the Statement Policy is legally correct as discussed in Section 2.

In the first sentence, we define *borrowing items*, *current*, *valid*, *identification card* and *issued by the University* or *issued by an affiliated program* as terms. The rationale behind this is that all of the terms, except *borrowing items*, are conditions under which a person is permitted to borrow a library item. Terms *issued by the University* and *issued by an affiliated program* are in disjunction since both types of cards can be used for borrowing library items. For this reason we interpret this sentence as a permission, since the person is only permitted to borrow books if the conditions mentioned above are met. The conditions are: having a card issued by the university or by an affiliated party program, which is valid, current and it can also serve as an ID card. This specific annotation can be found in the legislation in the demo account which is entitled "AUP Library Policy".

Sentence 2. contains just two important facts: *borrowing* and *personal use*. We interpret this sentence in a negative way as a prohibition to borrow library items in the case of non-personal use.

Sentence 3. contains two terms as well: *AUP ID card* and *personal use*. Although the meaning of this statement is vague (for example it uses only the term „*personal*“ and not „*personal use*“ as in the previous sentence) we deduce that it means that an AUP issued card shall be used only personally. We interpret this sentence as an obligation, since having a card issued by AUP is a necessary condition for personal use. The rest of the Policy Statement and test cases shall be formalized in the same manner, however other sentences do not contain contradictory statements and therefore these won't be further discussed and demonstrated in this paper. Their full annotation is still available in the demo account.

Given these three statements we return to the original problem, which was introduced in Section 2.: a person having a current and valid card issued by an affiliated party lends the card to another person. This person, who is also in possession of a special access card, borrows a book from the library. If we apply some reasoning over these facts based on our interpretation, then we shall have the following assumptions: we have a current and valid card issued by an affiliated program as well as a special access card, which derives also an ID card (sentence 5.). That means we are permitted to borrow, because we meet all the conditions in sentence 1. However, it is our friend who uses the card so we must also assume a non-personal use. According to this, it is prohibited to borrow according to sentence 2.

We can now ask the tool whether it is permitted to borrow or whether it is prohibited to borrow. The NAI infers that given the assumptions of the case, a person is permitted to borrow, because all the conditions are met. At the same time, the tool can also deduce that the person is prohibited to borrow. Under these circumstances, the tool is not only able to answer the question of whether a person is permitted to borrow but also to state the existence of an inconsistency in the policy.

³ Finding a correct formalization of a legal text is a task itself for which we use a behaviour-driven development methodology. The description of this methodology and its application for using the NAI tool is a future work for authors.

4. Finding and correcting inconsistencies in legal texts

In the previous section we have shown how two contrary conclusions can be derived from the AUP library regulations under a certain interpretation. We have followed the legal argument with a formalization in the NAI tool. A legal formalization from which we can derive such a contradiction is logically inconsistent. The NAI tool supports the automatic check of the consistency of a formalization by clicking the “*Run consistency check*” button. A consistent formalization has at least one interpretation in which all sentences are true.

Does it remain consistent when we assume further facts? The answer is no. Let us take, for example, the following two norms:

1. *If the end of the world has come, then you are permitted to smoke a cigarette.*
2. *If the end of the world has come, then you are forbidden to smoke a cigarette.*

The above two norms are together logically consistent. One interpretation in which all sentences are true is that in which the end of the world has not come and you are permitted to smoke a cigarette. If we further assume that the end of the world has come, then this interpretation is no longer possible. In fact, the two norms together with the additional fact are logically inconsistent, since in any possible interpretation, one of the sentences must be false.

While we might be fine to assume that the fictitious fact from above can never materialize and therefore that these two norms are consistent, this is not the case in general. If a consistent legal formalization contains some scenarios, which, by assuming them, can lead to inconsistencies, then we would like to consider the legal formalization as problematic. In other words, while the formalization is consistent and can be used to derive some correct conclusions, some other, incorrect, conclusions can be obtained when assuming some plausible facts. We are therefore interested in being able to detect when legal formalizations are inconsistent according to this wider notion.

4.1. Proving the inconsistency of legal formalizations

As mentioned in the previous section, clicking the “*Run consistency check*” button on the legislation editor only assures us that given no further facts, the legislation is consistent. It can become inconsistent if some plausible facts are assumed. The approach to checking for inconsistent formalizations is then to assume different facts and use the similar “*Run consistency check*” button which appears in the query editor. For this approach to work, we need to try different combinations of plausible facts. In the formalization described in Section 3., the case where a student of an affiliated program is giving her card to a friend for borrowing books from the Library triggers an inconsistency. This case is represented by the following facts: *current*, *valid*, *issued by an affiliated program*, *special access card* and *NOT personal use*.

How can we find which cases trigger inconsistency? In [LIBAL/STEEN 2019] a methodology is discussed which can be used in order to validate the correctness of formalizations. This methodology depends on having a good coverage of plausible cases - the user who follows this methodology should create different “*Test*” queries which describe different plausible scenarios. Correctness of the formalization is then being validated by checking that all queries are validating. What happens when a particular “*Test*” query describes a case which triggers inconsistency? In this case, we just need to click the “*Run consistency check*” button on this “*Test*” query and we know immediately that the formalization can become inconsistent. As mentioned in

Section 3, the prover not only tells us that this case triggers inconsistency but also provides us with a mathematical proof to this effect.

One should also mention here that this part of the methodology can be fully automated. This direction is described in more detail in Section 5.

4.2. Correcting inconsistent legal texts

Once we have obtained a case which triggers inconsistency and since we are assured that no error in the reasoning process can happen, we are faced with two possible scenarios. Either, our formalization of the legal text is incorrect or it is indeed the case that the legal text can introduce inconsistencies. As discussed in the previous sections, we can greatly decrease the chance of an incorrect formalization (but not eliminate it completely) by following the methodology in [LIBAL/STEEN 2019]. We will therefore assume that the legal text can introduce inconsistencies and describe a method for eliminating it.

Logically speaking, the source of the inconsistency lies in the existence of one or more norms which can be used to deduce contradictory conclusions. In our formalization from Section 3., the norm in sentence 2. triggers inconsistency in certain cases. This is so since there is vagueness about what personal use means in this case and about the precise context and priority of this norm, as discussed in Section 2.

There are multiple ways one can change the text in order to exclude this specific inconsistency. One can follow the *lex specialis* approach and merge sentences 2. and 3. On the other hand, one can remove the confusing sentence 3. Both are possible using the following methodology.

In the general case, the user needs to firstly identify the conflicting norms for a specific case. These norms need then to be corrected in a way that the inconsistency can no longer be triggered. According to the methodology discussed in the Section 4., the user obtains a set of “*Test*” queries which are used to validate the correctness of the formalization. In the case we are considering, at least one of these queries turns out to trigger inconsistency, as is validated by running the consistency check. For each query which triggers inconsistency, we suggest the following steps.

Identification of conflicting norms. The first step is to identify all norms which might cause the inconsistency. This step is currently done in a semi-automated way only. Section 5 discusses how this process can be fully automated. In order to identify the norms, we create new queries whose set of assumptions are identical to the assumptions of the inconsistent “*Test*” query. The goal of the queries are the assumptions of each norm. For example, if a norm assumes conditions A AND B AND C, the query testing this norm has as a goal A AND B AND C.

In our example, we have defined 8 norms and we therefore need to use 8 queries in order to identify the relevant ones. The NAI editor allows copying and pasting of already annotated text, which facilitate this step. By applying this step to the 8 norms, we obtain that the relevant ones are 1, 2 and 3.

Resolving inconsistencies. In this step, the user focuses on the relevant norms and attempts to resolve the conflict. Every such attempt is accompanied by running all “*Test*” queries. The goal of this step is to get the inconsistent query to become consistent while still having all other “*Test*” queries validating. Currently, this step is done manually. Section 5. considers some possibilities for automation.

In our example, we have merged sentences 1. and 2. and deleted sentence 3. altogether. This means that we have added the condition that the card is personal to the first norm, obtaining:

*“In order to borrow library items (books, films, CDs, etc.) the user must present a **personal**, current and valid identification card issued by the University (or an affiliated program)”*.

We proceeded by executing all “*Test*” queries. They all validate and the inconsistent case is no longer inconsistent. Although the ID card can be considered to be personal intuitively, on the other hand person can use cards issued by different affiliated programs with different conditions and adding “personal” unifies the situation for all cards.

5. Conclusion and future work

Inconsistencies in legal texts are sometimes hard to find since they may only be triggered in certain cases. In this paper we have described a methodology for the detection and correction of such inconsistencies and demonstrated it on a relatively simple case. However, legal practice can bring us more complexed examples of legal inconsistencies.

There are both technical and applicative possible extensions to the work discussed in this paper. The example chosen to illustrate our methodology in this paper is relatively simple and we would like to consider more interesting inconsistencies in legal texts in the future.

On the technical side, the methodology defined in this paper is based on the existing functionality of the NAI tool. There are several ways in which improvements in the tool can greatly enhance the methodology described in this paper. Currently, consistency checks can be executed for specific cases only. This limitation increases the manual work required in order to detect inconsistencies as the user must create various cases. Clearly, given a propositional formalization, there can be only a finite number of such cases, which can be enumerated according to the terms used. Even if first-order logic is used for formalizing the text, finite models can still be found and used in order to automate cases generation. We plan on integrating this functionality as a synchronous operation in a future version of NAI.

A similar automation can also be obtained for the detection of the formulae which are part of the inconsistency. The tool needs only to try to infer the conditions of each formula from the facts of the inconsistent case. One can also argue further that the text can also be corrected automatically, at least in some cases. This can be achieved by trying different combinations of the problematic formulae while running all cases at the same time. A correct formalization is obtained when all cases validated are consistent.

References

- BAR-HAIM, ROY et al., The Second PASCAL Recognising Textual Entailment Challenge, In: Proceedings of the Second PASCAL Challenges Workshop on Recognising Textual Entailment, 2006.
- CLARK, PETER et al., An Inference-Based Approach to Recognizing Entailment, In: Proceedings of the Fifth PASCAL Challenges Workshop on Recognising Textual Entailment (RTE 2009), 2009.
- CONDORAVDI, CLEO et al., Entailment, Intensionality and Text Understanding, In: Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies 2003 (NAACL-HLT 2003), USA: Association for Computational Linguistics, 2003.
- COOPER, ROBIN et al., Using the Framework, 1996.
- DAGAN, IDO et al., Recognizing Textual Entailment. Rational, Evaluation and Approaches, In: Nat. Lang. Eng., 2009.
- DWORKIN, RONALD, No Right Answer?, New York University Law Review, 1978.

GIAMPICCOLO, DANILO et al., The Third PASCAL Recognizing Textual Entailment Challenge, In: Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing (RTE 2007), 2007.

HARABAGIU, SANDA M. et al., Negation, Contrast and Contradiction in Text Processing, In: Proceedings of the Twenty-First National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference (AAAI 2006), 2006.

KARLOVA-BOURBONUS, NATALI, Automatic detection of contradictions in texts, 2018.

KRESS, GUNTHER R., Linguistic processes in sociocultural practice, B.m.: Oxford University Press, 1989.

LIBAL, TOMER/STEEN, ALEXANDER, The NAI Suite - Drafting and Reasoning over Legal Texts. In the proceedings of JURIX, 2019. To appear.

MACCARTNEY, BILL et al., An Extended Model of Natural Logic, In: Proceedings of the Eighth International Conference on Computational Semantics (IWCS-8), 2009.

TATU, MARTA et al., COGEX at RTE3, In: Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing (RTE 2007), 2007.

TUSHNET, MARK V., Critical Legal Theory, The Blackwell Guide to the Philosophy of Law and Legal Theory. B.m.: Blackwell, 2005.