

## ANNOTATION OF RHETORICAL ROLES AND SYLLOGISTIC RELATIONS IN CZECH ARGUMENTATIVE LEGAL AND ADMINISTRATIVE TEXTS

SILVIE CINKOVÁ<sup>1</sup> – JANA ŠAMÁNKOVÁ<sup>2</sup> – BARBORA KUBÍKOVÁ<sup>3</sup>  
– TEREZA NOVOTNÁ<sup>4</sup> – VÍTEK EICHLER<sup>5</sup>

<sup>1</sup>Institute of Formal and Applied Linguistics, Faculty of Mathematics and Physics,  
Charles University, Prague, Czech Republic (ORCID: 0000-0003-4526-3915)

<sup>2</sup>Department of Legal Skills, Faculty of Law, Charles University, Prague, Czech  
Republic (ORCID: 0009-0008-6251-8223)

<sup>3</sup>Office of the Public Defender of Rights and Defender of Children's Rights, Brno,  
Czech Republic (ORCID: 0009-0002-7226-389X)

<sup>4</sup>Institute of Formal and Applied Linguistics, Faculty of Mathematics and Physics,  
Charles University, Prague, Czech Republic (ORCID: 0000-0002-1426-4547)

<sup>5</sup>Faculty of Law, Charles University, Prague, Czech Republic  
(ORCID: 0009-0008-3070-8997)

CINKOVÁ, Silvie – ŠAMÁNKOVÁ, Jana – KUBÍKOVÁ, Barbora – NOVOTNÁ,  
Tereza – EICHLER, Vítěk: Annotation of Rhetorical Roles and Syllogistic Relations in  
Czech Argumentative Legal and Administrative Texts. *Journal of Linguistics*, 2025, Vol. 76,  
No 1, pp. 145 – 156.

**Abstract:** We present KUKY 1.0, a publicly available corpus annotated by lawyers  
for rhetorical structure and relevance. We explain the concepts as well as the annotation  
procedure. Eventually we discuss the challenge of inter-annotator agreement.

**Keywords:** rhetorical structure, argumentation, annotation, corpus, legal, law,  
syllogism, relevance

## 1 INTRODUCTION

### 1.1 Readability of legal and administrative documents

This study describes a newly released annotated corpus of Czech legal and  
administrative documents with readability assessment and a multi-layer annotation  
of relevance, rhetorical roles, and, partly, of rhetorical structure. The purpose of the  
corpus is prototyping readable Czech legal/administrative documents with large  
language models.

There are enough grammatical and lexical strategies to increase readability,  
such as shorter sentences, avoiding passives, or replacing deverbal names with  
verbs (DuBay 2004) – strategies that work across languages. Some of the classic  
readability formulas have even been adapted to Czech (Bendová and Cinková  
2021). So it seems all what legal writers should do is to internalize these rules

and check results with a formula. However, recent research suggests that the effect of grammatical and lexical features on readability is rather modest (Cinková 2024).

This may not surprise lawyers, who have anyway traditionally concentrated on rhetorical structure as a component of logical coherence rather than on stylistics *per se*. Therefore, we have shifted focus in this very direction, and we involved lawyers to shape the annotation scheme. The resulting annotation scheme faithfully implements the steps plain legal writing experts take when perusing a document to optimize it for readability.

## 1.2 Related work

In the last decades, annotated legal corpora were produced to aid automatic *summarization* and *argumentation mining*. Their annotation schemes either capture the conventional macrostructure, such as preambles and decisions sentences (de Vargas Feijó and Moreira 2018; Šavelka and Ashley 2018), or they concentrate on the pragmatics of sentences or clauses.

For instance, Grover et al. (2003) analyzed judgments of the British House of Lords, using three labels: *Background* (references to law and precedents), *Case* (events and lower court decisions), and *Own* (speaker’s judgments and interpretations of *Background*). Bhattacharya et al. (2023), as well as Malik et al. (2022) arrived at a more fine-grained set of *Facts*, *Ruling by Lower Court*, *Argument* (of the present Court), *Statute* (laws references by the present court), *Precedent*, *Ratio of the Decision*, and *Ruling by Present Court*. The best-known legal corpus with argumentation mining annotation is the ECHR judgments corpus (Teruel et al. 2018), along with its recent extensions (Habernal et al. 2024). Yamada et al. (2019) built a corpus of legal argumentation of Japanese civil law judgments. Unlike in most corpora, the annotation segments are not strictly defined by single sentences or clauses but are allowed to span across. This annotation scheme provides both labels of rhetorical roles and relations between them.

KUKY 1.0 resembles the Japanese corpus by allowing for a free segmentation of annotation spans. The labels draw on Grover et al. (2003) and seek to link the corresponding spans into syllogistic triples of premises and conclusions. Apart from the rhetorical roles, the corpus provides an annotation of relevance. This annotation highlights incomprehensible or superfluous text.

# 2 RELEVANCE AND RHETORICAL ROLES

## 2.1 Relevance

Irrelevant information increases the cognitive burden of the reader in two ways: for the first, a longer message takes a longer time to read; for the second, the reader will waste their cognitive capacity on integrating disparate inputs to make sense as

a whole. Empirical research (Tyler 1990; Song and Schwarz 2010; Wagner and Walker 2019) has proven that the easier a message is to perceive, the more persuasive or authoritative it appears to the reader. Hence, irrelevant information hampers real-world processes in administration and justice by obscuring the actual messages.

When annotating relevance in KUKY 1.0, the annotators mimic the first step in redesigning documents, where the editor deliberates which original content to preserve in the new version. They mark spans as *Relevant*, *Superfluous/Irrelevant*, and *Incomprehensible/Confusing beyond repair*.

## 2.2 Syllogism in argumentative writing

Argumentative texts judge whether or not a fact contradicts the law. That requires setting out the relevant law and project it on a fact in such a way that the applicability of the given law to the given fact becomes indisputable. The stronger the link between the law and the fact, the more persuasive is the resulting judgment.

According to J. Gardner (1993), the rhetorical centerpiece of legal persuasive reasoning is *syllogism*<sup>1</sup>. Gardner argues that “all legal argument should be in the form of syllogisms” because “syllogistic argument provides the requisite appearance of certainty. It makes the outcome of a case seem as certain and as mechanical as the output of a mechanical equation, and achieves this effect not by actual mathematical operations, but, paradoxically, by exploiting human intuition” (Gardner 1993 §1.1.). Syllogism consists of three components:

1. The *major premise* (a broad statement of general applicability),
2. The *minor premise* (a narrower statement able to serve as an instance of the major premise),
3. The *conclusion* (a statement that evidently holds for the major premise, and, hence, it must also hold for the minor premise).

In the legal domain, the major premise is populated by the law, the minor premise by the fact, and the conclusion by the judgment. Hence, it is not by chance that the law, fact, and conclusion are the central labels in the annotation of rhetorical roles in legal corpora, and KUKY 1.0 is no exception.

## 2.3 Rhetorical roles in KUKY 1.0

KUKY 1.0 distinguishes between argumentative and normative documents. Tab. 1 presents the annotation scheme of argumentative documents.

Narrative	Minor premise. Facts, testimonies, and past decisions by authorities.
Law	Major premise. Law references, quotes, interpretations, and summaries.
Conclusion	Conclusion. Ruling, finding, judgment.
Advice	Optional information to aid the recipient.

<sup>1</sup> The concept of syllogism is attributed to Aristotle (Aristotle 2004).

Command	Recipient's obligations resulting from the document.
Legal Issue	Summary of the matter of dispute in legal terms, typically formulated as a yes-no question.
Metatext	Processing matters.

Tab. 1. Rhetorical roles in argumentative documents

## 2.4 Syllogism in KUKY 1.0

Among the argumentative documents in KUKY 1.0, the presence of syllogistic structures distinguishes the top-readable documents from the ordinary ones. We will illustrate syllogism on two authentic examples in Tab. 2 and Tab. 3.

Tab. 2 presents numbered and labeled segments of a court order. Court orders always start with the court ruling. A good court ruling is the summary of a Conclusion in the reasoning part. A good reasoning part contains a Conclusion, which can be divided into several partial ones, but each Conclusion must be backed up by at least one Law and one Narrative, which would ideally match each other in the syllogistic way.

The court uses syllogistic argumentation in this order. It rules that a certain Pavel Boháč can legally represent his elderly mother. The conclusions (1 and 5) are in accordance and Sentence 1 does not add anything new to 5. Conclusion 5 is supported by one Law (2) and one Narrative (3). Hence Conclusion 1 is supported by the same Law and Narrative. The Law and the Narrative match point by point: the familial relationships, the mother's health conditions, and her deliberate approval.

ID	Text span	Label
1	<p><b>Order</b>  <i>The District Court [...], has decided in the legal matter concerning Jitka Boháčová's approval for representation by a household member as follows:</i>  <i>The court approves the representation of Jitka Boháčová by her household member, Mr. Pavel Boháč.</i></p>	Conclusion
2	<p><b>Legal Framework</b>  <i>If a mental disorder prevents an adult from legally acting on their own behalf, they can be represented by a household member [...]. The representative must inform the represented [...] and clearly explain [...]. If the person to be represented refuses, the representation does not arise [...].</i>  <i>Court approval is required [...]. Before issuing a decision, the court must make the necessary efforts to ascertain the opinion of the represented person [...]</i></p>	Law
3	<p><b>Assessment of the petition</b>  <i>The court verified that there is indeed a familial relationship [...]. The court visited the subject [...]. The court verified that the subject's health condition [...]. The court confirmed that the subject understands the nature and consequences of the representation, agrees with it, and agrees that the petitioner, her son, will represent her.</i></p>	Narrative

4	<i>The court thus found that the conditions for approving the petitioner as the subject's representative were met [...]. The court also confirmed that representation by a household member is sufficient to protect the rights and interests of the subject.</i>	Conclusion
5	<b><i>Some Rights and Duties of the Household Member Representative</i></b>	Advice

**Tab. 2.** A court order about legal representation by a household member

Another example (Tab. 3) shows the use of syllogism in a last-warning letter for a neighbour to confine her overgrown trees to her lot. Note how two rhetorical roles can appear within one sentence: the first sentence starts with a Conclusion (1) and continues with a Narrative (2). In the Czech original, the narrative is structured as a subordinate content clause (*tím, že...*). The supporting Law appears in 3 and 5. In addition, two Commands (4, 6) are each backed up by a Law (3, 5) as well as by the Narrative, and act very much as Conclusions and actually form two other syllogisms. So, virtually all statements in this documents are components of a syllogism, and this is what makes the text particularly succinct and the train of thought so easy to follow.

ID	Text span	Label
1	<i>I am informing you that you are violating my property rights.</i>	Conclusion
2	<i>by having planted and grown trees in close proximity of the border between our lots as well as continuously planting new trees without maintaining them, so that their branches and roots are reaching over to my lot.</i>	Narrative
3	<i>According to §1016 of the Civil Code you are obliged to maintain all hanging and underground parts of your trees that trespass the border on my lot.</i>	Law
4	<i>Therefore I urge you to cut the branches that reach over on my lot and to remove the undergrowing roots on my lot, all of this within 30 days from the delivery of this letter.</i>	Command
5	<i>According to §1017 of the Civil Code you can only plant tree species that usually grow above 3 meters at least 3 meters from the border of your lot, lower growing trees then at least 1,5 meters from the border of your lot.</i>	Law
6	<i>Therefore I urge you to stop planting new trees at our common lot border in a way that contradicts the law.</i>	Command
7	<i>I firmly believe that you are going to stand up to your obligations. Otherwise I will take you to court.</i>	Advice

**Tab. 3.** Syllogistic structure in a final demand

### 3 DATA

#### 3.1 Statistics

KUKY 1.0 is a curated selection of 224 Czech administrative and legal documents (totalling of 374,251 tokens) for readability research, formatted in plain

text with or without markdown. The document length lies between 159 and 6,239 tokens, with the median at 1,250 and the mean at 1,671 tokens.

Document contributors were legal experts dedicated to plain legal writing, who sought to select a range of examples from high-quality documents (to serve as blueprints for the given genre), over somewhat accessible documents, to the standard production, which is generally hard to comprehend. The collection is somewhat biased towards the best and good documents, since they require a more careful selection than the standard production, which can be acquired bulk-wise from other sources.

### 3.2 Document sources

The main sources of documents in KUKY 1.0 are the publicly available databases of the Office of the Czech Public Defender of Rights, the Supreme Administrative Court, and a free legal advice database of a legal company (Frank Bold). Besides, the corpus contains various contributions from individual legal experts: communications between clients and authorities, public local administration announcements, or legal memos. Such documents were thoroughly pseudonymized, including local names, dates, and all other numeric strings, to preclude tracing of the parties involved.

On the top level, the documents are grouped into *argumentative* (174 documents) and *normative* (50 documents). Argumentative documents are always case-related. They map a past event or its result on existing legal norms to judge it. Typical argumentative documents are findings and decisions by authorities, such as courts and supervisory bodies, or personalized client advice by legal experts. Normative documents, on the other hand, set norms (laws) or guide a generic reader through an administrative procedure (e.g. how to register a society). They can even model life situations (e.g. how to deal with a noisy neighbor), but they never address a concrete case.

### 3.3 Metadata and structure

The main document distinction is the *argumentative* vs. *normative*, but a few more criteria were used to classify the documents and captured in the metadata by single judgments of document contributors (Tab. 4).

The argumentative documents and the normative documents come in two JSON files. Both files consist of three JSON arrays: *documents*, *labels*, and *annotations*. The *documents* array contains objects that represent the individual documents. Each object in the *documents* array contains the document's text, along with metadata, as object properties. The *labels* array lists the labels defined by the annotation scheme. The annotation schemes of the argumentative and the normative documents slightly differ, which is why they are stored in separate files. The *annotations* array lists individual annotations: texts spans marked with labels. Each object within the *annotations* array contains an annotation label and maps on the source text with a reference to the document's ID and with offsets.

Each document has two annotation layers: the rhetorical roles and relevance. These two layers were annotated independently, so each segments the text differently. Their definitions in the *labels* JSON object are merged, but their instances in the *annotations* objects are distinguished by a property called *task\_type*.

### 3.4 Access

The entire KUKY 1.0 corpus along with the documentation is stored in the LINDAT/CLARIAH-CZ repository under the persistent ID <http://hdl.handle.net/11234/1-5812> and a CC BY-NC-SA 4.0 license. The documentation is also available at the non-persistent URL <https://ufal.mff.cuni.cz/grants/ponk/kuky>.

Metadata property	Values	Description
<i>doc_id, doc_name</i>		Unique document ID, name
<i>Readability</i>	<i>Low, Medium, High</i>	Expert assessment, relative to other documents in the corpus.
<i>Anonymized</i>	<i>Anonymized by source, On-site anonymization, No</i>	Is the document anonymized/pseudonymized?
<i>SyllogismBased</i>	<i>True, False</i>	Does this document systematically use syllogism?
<i>DocumentVersion</i>	<i>Original, Partial Redesign, Redesign</i>	Default: Original. Some documents come in an original version and revision(s).
<i>ParentDocumentID</i>		Redesigned documents contain a reference to the <i>doc_id</i> of their corresponding Original.
<i>LegalActType</i>	<i>Individual, Normative</i>	The key distinction between documents in this corpus.
<i>Objectivity</i>	<i>Quasiobjective, Persuasive</i>	Judgments are quasiobjective. Lawsuits etc. are persuasive.
<i>Bindingness</i>	<i>True, False</i>	Is the document legally binding?
<i>AuthorType</i>	<i>Authority, Individual</i>	Does the author write in the capacity of an authority?
<i>RecipientType</i>	<i>Natural person, Legal person, Combined</i>	Natural persons are not likely to hire an expert to interpret the document for them, while legal persons (e.g. companies) often employ lawyers.
<i>RecipientIndividuation</i>	<i>Individual, Bulk, Public</i>	How familiar are the recipients with the matter?

Tab. 4. Metadata properties in documents

## 4 ANNOTATION

### 4.1 Procedure

The annotation proceeded in two separate steps: relevance and rhetorical roles including the syllogistic relations. Both steps were carried out in a cloud installation of Gloss (Poudyal et al. 2020), by courtesy of its developer Jaromír Šavelka. The texts were selected, assessed, edited, and subsequently annotated by lawyers, mostly ones with an extensive experience with practicing as well as teaching plain legal writing.

Deliberately segmented data pose a challenge for measuring the inter-annotator agreement (IAA). Differences in segmenting should not be penalized, as long as the words were identically labeled. Therefore we considered each token one annotator judgment. We report IAA for ten documents and two annotators.

### 4.2 Inter-annotator agreement on relevance

IAA on Relevance reached accuracy 0.78. Cohen's Kappa over all labels was only 0.47, which is not too bad considering that the *Relevant* label very strongly prevailed, and hence each disagreement was heavily penalized. The prevalence of the *Relevant* label is evident from the confusion matrix in Tab. 5.

	Incomprehensible	Superfluous	Relevant
Incomprehensible	183	495	196
Superfluous	0	999	2071
Relevant	128	437	10330

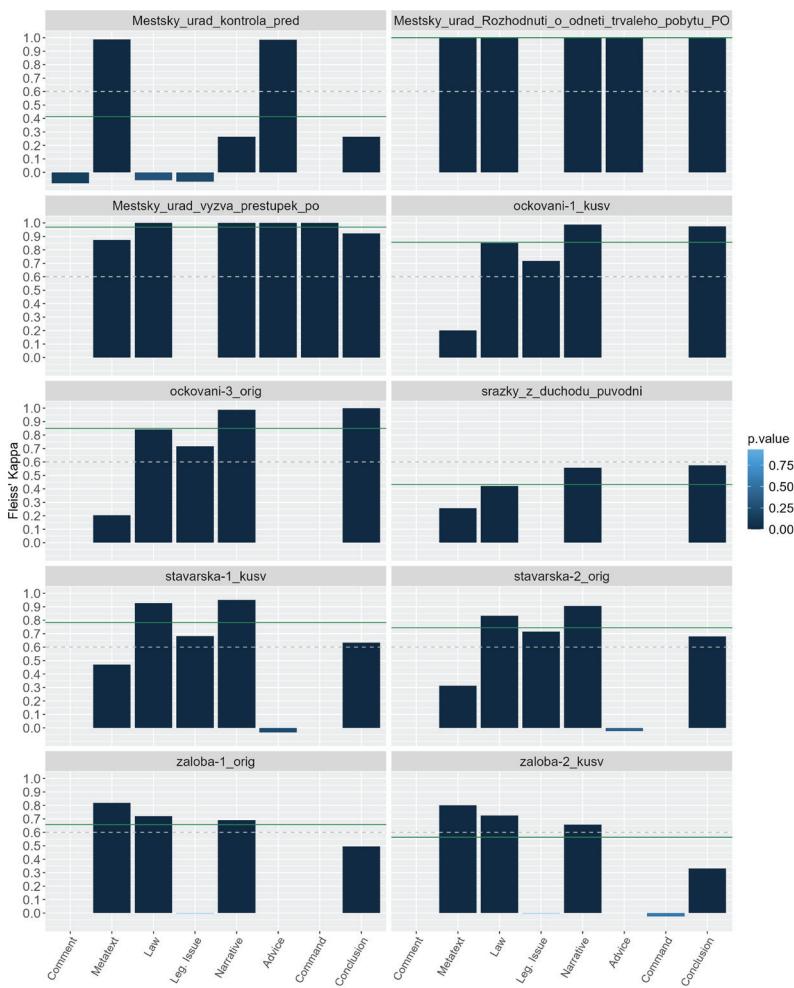
**Tab. 5.** Confusion matrix of Relevance annotation (numbers stand for count of tokens with the given combination of labels)

### 4.3 Inter-annotator agreement on rhetorical roles

IAA on rhetorical roles varied very strongly across documents. Fig. 1 illustrates the IAA as Fleiss' Kappa on individual labels within individual documents. The dashed line represents the average Fleiss Kappa across all labels within the document. The solid line is placed at 0.6, a rule-of-thumb threshold for semantic tasks.

### 4.4 Inter-annotator agreement on syllogistic relations.

To compute IAA on syllogisms, we modeled the relations between segments as relations between individual tokens. There were possible relations per document (each token with each token). Actual relations between segments were modeled on each word of one segment to each word of the second segment. We neglected their rhetorical role labels. The average accuracy was 0.95 (standard deviation 0.05), precision 0.53 (standard deviation 0.29), and recall 0.3 (standard deviation 0.31).



**Fig. 1.** Inter-annotator agreement on rhetorical roles, document-wise and label-wise

## 5 DISCUSSION

The tagsets mimic the deliberation phase of a human editor, pursued manually with crayons before drafting the redesigned version—from scratch, with occasional copy-pasting. Even though the two editors follow the same principles, they might pursue them differently, just as the resulting redesigns would never be identical across authors, although both could be equally good.

The IAA is not impressive, but this could be expected with a task that is closer to translation rather than classification. The Relevance annotation has a low IAA because the distribution of labels is very uneven. In practice, annotators recognize most of document content as relevant, no matter how clumsy the style: hence the corresponding accuracy of almost 80%.

The average IAA on Rhetorical Roles does not say much because of the wide dispersion among documents. In fact, only three documents of ten do not reach the 0.6 average Fleiss' Kappa. None of these three had been classified as highly readable in the metadata (before the annotation). Even a most shallow disagreement analysis reveals that rhetorical roles are blurred in unreadable documents, suggesting that the speaker does not care to organize their utterance into purposeful units.

In the document with the worst IAA (Mestsky\_urad\_kontrola\_pred), Narrative, Law, and Conclusion mingle with Metatext even within one sentence, such as in this example:

*V průběhu kontroly bylo zjištěno podezření z porušení ustanovení § 21 odst. 4 zákona o ochraně veřejného zdraví, kterého se kontrolovaná osoba dopustila tím, že v průběhu kontroly nebyl v kontrolované provozovně vyvěšen provozní řád schválený orgánem ochrany veřejného zdraví, a to v souladu s výše uvedeným ustanovením, přestože je v dotčené provozovně vykonávána činnost „Pedičura, manikúra“, která je zákonem o ochraně veřejného zdraví považovaná za činnost epidemiologicky závažnou.*

‘During the inspection, a suspicion of violation of instruction § 21 Par 4 of the Public Health Protection Law was detected, that the inspected person committed by the fact that during the inspection at the inspected shop the operation rules approved by the Public Health Protection officer were not on display, that in accordance with the aforementioned instruction, although in the aforementioned shop was carried out the activity “Pedicure, Manicure”, which is considered an epidemiologically relevant activity by the Public Health Protection Law.’

It goes without saying that IAA is hard to maintain when untangling such a scramble into discrete communicative intents. So, for instance, a poorly referenced law might be recognized as such by one annotator, while the other would “downgrade” it to Metatext. The same could easily happen to a sloppily formulated Legal Issue or a nebulous Conclusion.

The syllogistic annotation heavily depends on the Rhetorical Roles annotations. When a text contains numerous ambiguous segments or other forms of incongruity, annotators are often reluctant to scour it for potential partial syllogisms.

Qualitative observations suggest that comprehensible documents are easier to agree on, reminding us of the proverbial Anna Karenina principle saying that all happy families are alike, while each unhappy family is unhappy in its own way. We speculate that, in a machine-learning setup, readability assessment it is not going to

be aided as much by the automatic classification of the rhetorical roles themselves as by the confidence levels of the predictions, and the same would apply to the detection of syllogistic structures.

## ACKNOWLEDGEMENTS

The research has been supported by the Technology Agency of the Czech Republic under the project TQ01000526 (9/2023-12/2025).

## References

Aristotle (2004). *Rhetoric*. New York: Dover Publications.

Bendová, K., and Cinková, S. (2021). Adaptation of Classic Readability Metrics to Czech. In 24<sup>th</sup> International Conference on Text, Speech and Dialogue. Cham, Switzerland: Springer, pp. 159–171.

Bhattacharya, P. et al. (2023). DeepRhole: Deep Learning for Rhetorical Role Labeling of Sentences in Legal Case Documents. *Artificial Intelligence and Law*, 31(1), pp. 53–90. Accessible at: <https://doi.org/10.1007/s10506-021-09304-5>.

Cinková, S. (2024). Linguistic Factors in the Readability of Czech Administrative and Legal Texts. In: Z. Bohušová – M. Dove (eds.): *To Understand Is to Be Free. Interdisciplinary Aspects of Comprehensibility and Understanding*. Vienna, Austria: Praesens Verlag, pp. 303–325.

DuBay, W. H. (2004). *The Principles of Readability*. Costa Mesa, California: Impact Information. Accessible at: [https://www.researchgate.net/publication/228965813\\_The\\_Principles\\_of\\_Readability](https://www.researchgate.net/publication/228965813_The_Principles_of_Readability).

Gardner, J. A. (1993). Legal Argument: The Structure and Language of Effective Advocacy. LexisNexis. Accessible at: <https://store.lexisnexis.com/en-us/legal-argument-the-structure-and-language-of-effective-advocacy-sku-us-ebook-03082-epub.html>.

Grover, C., Hachey, B., and Korycinski, C. (2003). Summarising Legal Texts: Sentential Tense and Argumentative Roles. In Proceedings of the HLT-NAACL 03 Text Summarization Workshop, pp. 33–40. Accessible at: <https://aclanthology.org/W03-0505/>.

Habernal, I. et al. (2024). Mining Legal Arguments in Court Decisions. *Artificial Intelligence and Law*, 32(3), pp. 1–38. Accessible at: <https://doi.org/10.1007/s10506-023-09361-y>.

Malik, V. et al. (2022). Semantic Segmentation of Legal Documents via Rhetorical Roles. In: N. Aletras et al. (eds.): *Proceedings of the Natural Legal Language Processing Workshop 2022*. Abu Dhabi, United Arab Emirates (Hybrid): Association for Computational Linguistics, pp. 153–171. Accessible at: <https://doi.org/10.18653/v1/2022.nllp-1.13>.

Poudyal, P. et al. (2020). ECHR: Legal Corpus for Argument Mining. In: E. Cabrio – S. Villata (eds.): *Proceedings of the 7<sup>th</sup> Workshop on Argument Mining*. Online: Association for Computational Linguistics, pp. 67–75. Accessible at: <https://aclanthology.org/2020.argmining-1.8/>.

Šavelka, J., and Ashley, K. D. (2018). Segmenting U.S. Court Decisions into Functional and Issue Specific Parts. In *Frontiers in Artificial Intelligence and Applications*. IOS Press. Accessible at: <https://doi.org/10.3233/978-1-61499-935-5-111>.

Song, H., and Schwarz, N. (2010). If It's Easy to Read, It's Easy to Do, Pretty, Good, and True. *Bulletin of the British Psychological Society*, 23(2), pp. 108–111.

Teruel, M. et al. (2018). Increasing Argument Annotation Reproducibility by Using Inter-Annotator Agreement to Improve Guidelines. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*.

Tyler, T. R. (1990). *Why People Obey the Law*. New Haven and London: Yale University Press.

de Vargas Feijó, D., and Moreira, V. P. (2018). RulingBR: A Summarization Dataset for Legal Texts. In: A. Villavicencio et al. (eds.): *Computational Processing of the Portuguese Language*. Cham: Springer International Publishing, pp. 255–264.

Wagner, W., and Walker, W. (2019). Incomprehensible!: A Study of How Our Legal System Encourages Incomprehensibility, Why It Matters, and What We Can Do About It. Cambridge Core. Cambridge: Cambridge University Press. Accessible at: <https://doi.org/10.1017/9781139051774>.

Yamada, H., Teufel, S., and Tokunaga, T. (2019). Building a Corpus of Legal Argumentation in Japanese Judgement Documents: Towards Structure-Based Summarisation. *Artificial Intelligence and Law*, 27(2), pp. 141–170. Accessible at: <https://doi.org/10.1007/s10506-019-09242-3>.