

Semantic Types for Computational Legal Reasoning: Propositional Connectives and Sentence Roles in the Veterans’ Claims Dataset

Vern R. Walker

Director, Research Laboratory for
Law, Logic and Technology
Maurice A. Deane School of Law
Hofstra University, New York
Vern.R.Walker@Hofstra.edu

Xiang Ni

Research Laboratory for
Law, Logic and Technology
Maurice A. Deane School of Law
Hofstra University, New York

Ji Hae Han

Research Laboratory for
Law, Logic and Technology
Maurice A. Deane School of Law
Hofstra University, New York

Kaneyasu Yoseda

Research Laboratory for
Law, Logic and Technology
Maurice A. Deane School of Law
Hofstra University, New York

ABSTRACT

This paper announces the creation and public availability of a dataset of annotated decisions adjudicating claims by military veterans for disability compensation in the United States. This is intended to initiate a collaborative, transparent approach to semantic analysis for argument mining from legal documents. The dataset is being used in the LUIMA argument-mining project. We address two major sub-tasks for making legal reasoning computable. First, we report the semantic types of propositional connective we use to extract information about legal rules from sentences in statutes, regulations, and appellate court decisions, and to represent those rules as integrated systems. Second, we report the semantic types of sentence role we use to extract and represent the fact-finding reasoning found in adjudicatory decisions, with the goal of identifying successful and unsuccessful patterns of evidentiary argument. For each type system, we provide explanations and examples. Thus, we hope to stimulate a shared effort to create diverse datasets in law, to empirically evolve optimal sets of semantic types for argument mining, and to refine protocols for accurately applying those types to texts.

CCS CONCEPTS

• **Information systems** → Decision support systems; Data mining; • **Applied computing** → Law;

KEYWORDS

Semantic data, inference role, propositional-connective type system, sentence-role type system, ontology, legal rule, argument mining, argumentation mining, computational argumentation, data

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee, provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

ICAIL '17, June 12 - 16, 2017, London, UK
Copyright is held by Vern R. Walker. Publication rights licensed to ACM.
ACM XXXXXX
<http://XXXXX>

validity, protocols.

ACM Reference format:

V.R. Walker, J.H. Han, X. Ni, and K. Yoseda. 2017. Semantic Types for Computational Legal Reasoning: Propositional Connectives and Sentence Roles in the Veterans’ Claims Dataset. In *Proceedings of the 16th International Conference on Artificial Intelligence and Law, London, UK, June 2017 (ICAIL '17)*, 10 pages.

1 INTRODUCTION

This paper announces the creation and public availability of a dataset of annotated decisions adjudicating claims by military veterans for disability compensation in the United States. We hope that this initiates a collaborative and transparent approach to semantic analysis oriented toward argument mining from legal documents. This dataset is currently being used in the LUIMA argument-mining project being conducted by Carnegie Mellon University, the University of Pittsburgh, and Hofstra University’s Research Laboratory for Law, Logic and Technology (LLT Lab) [5, 6, 7, 12, 32]. This paper reports the semantic types that we have found useful in representing the legal rules extracted from the statutes, regulations, and appellate decisions, and in representing the evidentiary reasoning extracted from fact-finding decisions of this dataset. We intend this paper to support several reasons why a public, collaborative research effort is desirable.

First, it is important to develop semantic types for computational representations empirically. It is inefficient and potentially misleading to start from highly developed, abstract models based on a few samples, and then try to link linguistic features of actual texts to the elements of those abstract models. The numerous examples of texts we provide throughout the paper show the great diversity of linguistic expression that we need to link to semantic types. We are convinced that we need to develop the optimal sets of semantic types from the bottom up.

Second, the development of protocols for manually assigning semantic types to linguistic features is a critical step for training annotators, for assuring accuracy of annotations, and for developing software to automate the annotation of new documents.

Protocols for manually producing semantic data are also critical for confirmation of coding by others, for replication of experimental results, and for use in building separate but compatible datasets. (By using the word “protocol,” we intend to imply both a detailed plan for implementing a procedure, as well as a set of conventions or criteria governing the generation of data.) Generating diverse sets of reliable and valid data is impossible without agreed methodologies for creating such data.

Third, annotated datasets should be made public, so that they become available for researchers and open to scrutiny. Open discussion of shared data is the only hope for evolving sets of semantic types that are adequate and optimal for argument mining.

2 PRIOR RELATED WORK

In this paper, we address two major sub-tasks for making legal reasoning computable. First, we report the semantic types of propositional connective we use to extract information about legal rules from sentences in statutes, regulations, and appellate court decisions, and to represent those rules as integrated systems. Second, we report the semantic types of sentence role we use to extract and represent the fact-finding reasoning found in adjudicatory decisions, with the goal of identifying successful and unsuccessful patterns of evidentiary argument.

The work of representing systems of legal rules and patterns of evidentiary reasoning in a computable form draws on many strands of prior related work. [15, 25, 28] Our emphasis here, however, is on systematic efforts to empirically match semantic types to legal texts with the goal of argument mining. [30] Of the research in this area, we have space to mention here only a number of recent developments.

First, there is extensive research on classifying statutory texts into semantic types, and on extracting legal rules from statutes and regulations. We discuss our approach to this task in Section 4 of this paper. Prior work on classification includes: classifying provisions of state statutes and regulations in the U.S. for functional attributes relevant to public health emergency preparedness [27, 13]; automatically classifying legislative texts into a hierarchy of legal topics [23]; automatically classifying paragraphs of legislative texts into provision types (e.g., definition, prohibition, penalty) [10]; and comparing machine learning to a pattern-based classifier in automatically classifying sentences in Dutch legislation into categories (e.g., definition, obligation, repeal) [17]. Prior work on extracting legal rules includes: using a linguistically-oriented, rule-based approach to identify and extract legal rules from a U.S. regulation [36]; using simple natural language processing to extract concepts, relations and definitions from written sources of law in the tax domain [35]; identifying four logical connectives used to list conditions within legislation in Hungary [18]; combining a top-down ontological approach with a bottom-up linguistic approach for knowledge modeling and learning, in a case study involving English and Italian directives, regulations and judicial decisions concerning EU consumer protection law [11]; and using parsed versions of sentences in legislative texts to automatically extract complex textual expressions and formulate legal rules [9]. What we are not aware of is a systematic, empirical investigation into the semantic types of logical connective needed to extract information about legal rules from the sentences of statutes, regulations, and judicial decisions. Moreover, those semantic

types should be useful in integrating such information into a computational representation of a rule system that can be used to govern fact-finding in adjudicatory proceedings.

Second, there is prior research on extracting instances of arguments or reasoning from the decisions of fact-finding tribunals, with the objective of identifying patterns of reasoning that are successful or unsuccessful. We discuss our approach to this task in Section 5 of this paper. Prior work on adjudicatory decisions includes: applying machine learning to annotate sentence types in vaccine-injury compensation decisions [5, 6, 7, 12]; assigning rhetorical roles to sentences in Indian court decisions [26]; extracting rulings from court decisions and retrieving prior relevant cases [16]; classifying sentences as argumentative in the Araucaria corpus, including newspapers and court reports [21]; using tree structures to represent complex argumentation [20]; automatically summarizing legal judgments of the U.K.’s House of Lords, in part by classifying the rhetorical status of sentences [14]; using statistical parsing as the input for computing quasi-logical forms as deep semantic interpretations of sentences in U.S. appellate court decisions [19]; and using narrative summaries of a case’s important facts to determine stereotypical fact patterns that strengthen or weaken a legal claim [4]. [2] provides one overview of statement types involved in legal reasoning in cases, from evidence to a verdict. The approach we describe in this paper is unusual in developing semantic types for adjudicatory decisions with a view toward integration with the representations of legal rule systems. Moreover, it is unusual that protocols for data generation are developed to enable validation and replication of the semantic data, or that the semantic data are publicly available for confirmation of accuracy and replication of experimental results.

3 THE DATASET OF VETERANS’ DISABILITY CLAIMS

This section describes the statutory and regulatory structure, as well as the adjudicatory process, from which the dataset of annotated decisions is drawn.

3.1 Statutory and Regulatory Framework

In the United States, benefits for veterans of the United States Uniformed Services are administered by the U.S. Department of Veterans Affairs (“VA”), and include disability compensation, educational assistance, and other benefits. [3] The enabling statutes are codified in the United States Code, primarily in Title 38 (e.g., 38 U.S.C., Chapter 11, on compensation for service-connected disability or death). The implementing regulations of the Department of Veterans Affairs are codified in the Code of Federal Regulations, primarily in Title 38 (e.g., 38 C.F.R. Part 3, concerning adjudication). Judicial decisions interpreting these statutory and regulatory provisions are discussed in Section 3.2.

3.2 The Process for Adjudicating Disability Claims

Individual claims for compensation for a disability (“disability claims”) usually originate at a VA Regional Office (“RO”) or at another local office across the country. [3, 22] If the claimant is dissatisfied with the decision of the RO, she may file an appeal to the Board of Veterans’ Appeals (“BVA”). Although the BVA is an administrative appellate body that is part of the VA, it has the

statutory authority to decide the facts of each case *de novo*. [22] The BVA must provide a written statement of the reasons or bases for its findings and conclusions, and that statement “must account for the evidence which [the BVA] finds to be persuasive or unpersuasive, analyze the credibility and probative value of all material evidence submitted by and on behalf of a claimant, and provide the reasons for its rejection of any such evidence.” *Caluza v. Brown*, 7 Vet.App. 498, 506 (1995), *aff’d*, 78 F.3d 604 (Fed. Cir. 1996). The BVA’s workload has increased dramatically in the past few decades, e.g., reaching 55,713 decisions in fiscal year 2015. [22; 8 p. 4] The vast majority (98%) of appeals considered by the BVA involve claims for disability compensation. [8 p. 1] Therefore, identifying and monitoring the patterns of factual reasoning within the decisions of the BVA present a significant challenge for automated argument mining.

Our veterans’ claims dataset currently contains annotated decisions of the BVA. Initially we are focusing attention on cases claiming disability for posttraumatic stress disorder (“PTSD”). PTSD is a mental health problem that some people develop after experiencing or witnessing a traumatic event, such as combat or sexual assault. Proving a service-connected PTSD disability can pose significant proof problems.

The veteran may appeal the BVA’s decision to the U.S. Court of Appeals for Veterans Claims (the “Veterans Court”). [22] This Court reviews the BVA’s factual findings using the “clearly erroneous” standard, the BVA’s interpretations of statutes and regulations using a *de novo* standard, and the BVA’s legal conclusions using a standard of “arbitrary, capricious, ... abuse of discretion, or otherwise not in accordance with law.” [22] The Court may issue non-precedential, single-judge decisions; precedential three-judge decisions; or full-court (*en banc*) decisions. The great majority of dispositions of appeals are by single-judge decisions (in Fiscal Year 2015, there were 1,313 single-judge dispositions, compared to 26 multi-judge panel dispositions, and 5 dispositions by the court *en banc* [29]).

Either the claimant or the VA may appeal a Veterans Court decision to the U.S. Court of Appeals for the Federal Circuit, and from that decision to the U.S. Supreme Court. The Federal Circuit may only review questions of law, such as a constitutional challenge, or the interpretation of a statute or regulation relied upon by the Veterans Court. [3, 22] It reviews such legal determinations under a *de novo* standard, but except for constitutional issues it “may not review any ‘challenge to a factual determination’ or any ‘challenge to a law or regulation as applied to the facts of a particular case.’” *Kalin v. Nicholson*, 172 Fed.Appx. 1000, 1002 (Fed.Cir. 2006). In the 12-month period ending March 31, 2015, 103 appeals from decisions by the U.S. Court of Appeals for Veterans Claims were terminated by the decisions of judges of the Federal Circuit (as contrasted with other types of case resolution). [1, Table B-8]

4 MAKING THE LEGAL RULES COMPUTABLE

Given the heavy case load of the BVA, automated argument mining is an important goal. Software that is even modestly successful at extracting arguments made and accepted or rejected in new cases could render the adjudicatory system as a whole more transparent, fairer across similar cases, and more efficient for litigants and for decision makers. We report in this paper on the semantic types we have developed for machine learning and

software programming for BVA decisions. This section of the paper discusses a computational representation of the governing legal rules. In addition, we provide numerous examples of legal texts to illustrate the importance of empirically investigating the meaning of individual sentences, when the goal is to integrate the provided information into a single computable rule system.

4.1 Representing the Legal Rules

For analyzing a fact-finding decision of the BVA, we begin by representing the legal rules stating the conditions under which the BVA must order compensation to be paid. Those legal rules are established by a combination of statute, regulation, and appellate case law. A legal rule can be represented as a set of propositions, one of which is the conclusion and the remaining propositions are the rule conditions [32, 20]. As shown in Figure 1, we represent a legal rule by placing the conclusion at the top of an indented list of its conditions, with each condition preceded by a symbol for

The veteran has a disability that is “service-connected.” [38 U.S.C.A. §§ 1110, 101(16) (2017); 38 C.F.R. § 3.303(a) (2017)]
&■ [1 of 3] The veteran has “a present disability.” [*Shedden v. Principi*, 381 F.3d 1163, 1166-67 (Fed.Cir. 2004); *Morris v. Shinseki*, 678 F.3d 1346, 1353 (Fed.Cir. 2012)]
V [1 of ...] The veteran has “a present disability” of posttraumatic stress disorder (PTSD), supported by “medical evidence diagnosing the condition in accordance with [38 C.F.R.] § 4.125(a).” [38 C.F.R. § 3.304(f) (2017)]
V [2 of ...] ...
&■ [2 of 3] The veteran incurred “a particular injury or disease ... coincident with service in the Armed Forces, or if preexisting such service, [it] was aggravated therein.” [38 C.F.R. § 3.303(a) (2017); *Shedden v. Principi*, 381 F.3d 1163, 1166-67 (Fed.Cir. 2004); *Morris v. Shinseki*, 678 F.3d 1346, 1353 (Fed.Cir. 2012)]
V [1 of ...] The veteran’s disability claim is for service connection of posttraumatic stress disorder (PTSD), and there is “credible supporting evidence that the claimed in-service stressor occurred.” [38 C.F.R. § 3.304(f) (2017)]
V [2 of ...] ...
&■ [3 of 3] There is “a causal relationship [“nexus”] between the present disability and the disease or injury incurred or aggravated during service.” [*Shedden v. Principi*, 381 F.3d 1163, 1166-67 (Fed.Cir. 2004); *Morris v. Shinseki*, 678 F.3d 1346, 1353 (Fed.Cir. 2012)]
V [1 of ...] The veteran’s disability claim is for service connection of posttraumatic stress disorder (PTSD), and there is “a link, established by medical evidence, between current symptoms and an in-service stressor.” [38 C.F.R. § 3.304(f) (2017)]
V [2 of ...] ...

Figure 1. Partial Rule Tree (List of Legal Rule Conditions) for Proving That a Veteran Has PTSD That Is Service-Connected

the logical connective operating between it and the conclusion. Each condition can function in turn as a conclusion, with its own conditions listed below it. The resulting nested sets of conditions has a tree structure – with the entire representation of legal rules being called a “rule tree” [32]. Figure 1 presents a partial rule tree that lists the primary rule conditions for proving that a veteran claimant has PTSD that is causally connected to a stressor (dangerous or traumatic event) that occurred during active service (i.e., that the veteran has service-connected PTSD). As will become clearer when we discuss propositional connectives in Section 4.2, Figure 1 shows that a claimant must prove a service-connected disability by proving all of three primary conditions (those preceded by “&■”); and if that disability happens to be PTSD, then there are specific conditions (preceded by “V”) for proving each of these three main conditions. As a result, in a BVA decision on a disability claim for PTSD, we expect the fact-finding reasoning to be oriented toward these three PTSD rule conditions.

For each propositional element in the rule tree, there is a citation to the authoritative source(s) of the rule, shown in square brackets after the proposition. The quotation marks for each proposition indicate which key words or phrases have been taken directly from these sources. Sources are primarily of three kinds: statutes (e.g., “38 U.S.C. § 1110”), administrative regulations (e.g., “38 C.F.R. § 3.304”), and judicial decisions that interpret the meaning of statutes or regulations (e.g., “*Shedden v. Principi*, 381 F.3d 1163 (Fed.Cir. 2004)”). An integrated system of legal rules (such as that displayed in Figure 1) typically combines rules from all three sources into a single computational system of rules.

4.2 The Semantics of Propositional Connectives in Legal Rules

Figure 1 illustrates how logical connectives link the propositional conditions of a legal rule to their conclusion. In Figure 1, the primary rule conclusion (at the top) is the proposition “The veteran has a disability that is ‘service-connected.’” In order to prove this proposition to be true, three primary conditions must be proved (listed and indented directly under the primary conclusion): (1) “The veteran has ‘a present disability’”; (2) “The veteran incurred ‘a particular injury or disease ... coincident with service in the Armed Forces, or if preexisting such service, [it] was aggravated therein’”; and (3) “There is ‘a causal relationship [‘nexus’] between the present disability and the disease or injury incurred or aggravated during service.’” This subsection of the paper discusses the logical connectives needed to represent the inferential relations from conditions to conclusion that we find in legal rule systems in the U.S., and thus needed to complete the partial rule tree shown in Figure 1. For example, in Figure 1, the inferential relation from the three primary conditions to the primary conclusion is, as a matter of law, a supportive, deductive inference from a set of jointly sufficient and independently necessary conditions. The phrase “as a matter of law” is used in U.S. law to designate that an inference is warranted by a pronouncement of a legal authority or source, even if not by common meanings, science, or common sense. It indicates that the inference is part of the formal semantics of the legal rule system, with the inferences being established by authoritative legal rules.

Inferential relations for legal rules have two basic properties or features: polarity and logical functionality. The definitions are

provided here, and examples are provided in the following discussion.

The **polarity** of an inferential relation is whether the truth of the condition(s) supports the conclusion (proves, or tends to prove, the conclusion to be true) or opposes the conclusion (proves, or tends to prove, the conclusion to be false).

The **functionality** of an inferential relation is whether the truth of the condition(s) either supports/opposes the conclusion in a deductive manner, or supports/opposes the conclusion in a probabilistic manner. If a rule is deductive, then it is straightforward to compute a truth-value for the conclusion given the truth-value(s) of the condition(s). (We do not comment here on which logic system would provide the best framework for formalizing deductive inference – while material implication and equivalence might be adequate, there may be reasons to prefer a many-valued, modal or deontic logic.) If a rule is probabilistic, then we might use probability models to compute a probability-value or truth-value for the conclusion. In our experience, the large majority of legal rules established by statute or regulation create deductive rules. Courts that are evolving legal rules are more likely to create probabilistic rules by establishing which factors or considerations, as a matter of law, are relevant to proving a conclusion.

Table 1 lists the eight propositional connectives that we have found to be necessary to represent the logic of legal rules. [Compare 30.] In the following subsections, we first explain the propositional connective and then provide examples of sentences whose representation employs that connective. We stress, however, that the empirical process of developing this typology of connectives followed the reverse process. We deployed new connectives only as needed to represent different types of inferential relations within stated rules. Each connective on the list is, of course, familiar from standard logics, but our emphasis here is on the development of an empirical basis for suggesting that the list as a whole is adequate to represent the information we have encountered thus far in legal texts.

Table 1. Eight Semantic Types of Propositional Connective

&	Jointly Sufficient Set of Conditions
&■	Jointly Sufficient Set of Necessary Conditions
V	Set of Independently Sufficient Conditions
■V	Necessary Set of Independently Sufficient Conditions
RF	Merely Relevant Condition
IRF	Irrelevant Condition
REBUT	Rebutting Defeater
UNDCT	Undercutting Defeater

4.2.1 A Jointly Sufficient Set of Conditions. If a set of conjoined conditions is deductively sufficient, as a matter of law, to prove the conclusion, then this is a **weak conjunction** (indicated by placing the symbol “&” before each condition of the set). The legal rule states that if all of the conditions in the set are true, then this is sufficient proof for establishing the truth of the conclusion, as a matter of law (i.e., the conclusion must be true, as a matter of legal authority).

For example, the statute establishes the basic entitlement to compensation for a disability with the following sentence:

For disability resulting from personal injury suffered or disease contracted in line of duty, or for aggravation of a preexisting injury suffered or disease contracted in line of duty, in the active military, naval, or air service, during a period of war, the United States will pay to any veteran thus disabled and who was discharged or released under conditions other than dishonorable from the period of service in which said injury or disease was incurred, or preexisting injury or disease was aggravated, compensation as provided in this subchapter,

38 U.S.C.A. § 1110 (2017) (Basic entitlement). This sentence establishes a legal rule with the conclusion “the United States will pay to [the] veteran . . . compensation as provided in this subchapter,” which is true as a matter of law under several conditions. The supporting conditions stated here are: “[the veteran has a] disability resulting from personal injury suffered or disease contracted in line of duty, or [from an] aggravation of a preexisting injury suffered or disease contracted in line of duty, in the active military, naval, or air service, during a period of war” and “[the veteran] was discharged or released under conditions other than dishonorable from the period of service in which said injury or disease was incurred, or preexisting injury or disease was aggravated.” Insofar as we are parsing this single sentence in isolation, we are warranted only to interpret it as establishing a weak conjunction of its supporting conditions: **if** all of the conditions are satisfied, **then** the United States will (must) pay compensation. This sentence by itself does not state that these are the only conditions under which the United States will pay compensation.

4.2.2 A Jointly Sufficient Set of Necessary Conditions. If a set of conjoined conditions is deductively sufficient, as a matter of law, to prove the conclusion, but the conclusion can be true only if every one of those conditions is true, then this is a **strong conjunction** (indicated by placing the symbols “&■” before each of the necessary conditions). The addition of the symbol for necessity (“■”) indicates that the condition is individually necessary for proving the conclusion. The legal rule states that, as a matter of law, not only do the conjoined conditions as a set provide sufficient proof for the truth of the conclusion (i.e., the conclusion must be true, as a matter of legal authority), but also the conclusion can be true only if each of the necessary conditions is true (the truth of the conclusion entails the truth of each of the necessary conditions).

We may encounter such strong conjunctions when the legal rule states the necessary elements for a valid legal claim (as in Figure 1), or when a rule creates an exception to a broader rule (i.e., the courts tend to interpret exceptions “narrowly,” as requiring that certain conditions must be satisfied). For example, in Figure 1, the three primary conditions for proving that “The veteran has a disability that is ‘service-connected’” are also interpreted by the court in *Shedden v. Principi*, 381 F.3d 1163 (Fed.Cir. 2004), as being necessary conditions. The decision states:

The Court of Appeals for Veterans Claims has correctly noted that in order to establish service connection or service-connected aggravation for a present disability the veteran must show: (1) the existence of a present disability; (2) in-service incurrence or aggravation of a disease or injury; and

(3) a causal relationship between the present disability and the disease or injury incurred or aggravated during service.

Shedden, 381 F.3d at 1166-67. The court lists three necessary conditions for proving service connection, by stating what the claimant “must show.” Thus, *Shedden* provides legal authority for regarding the conjoined three main conditions in Figure 1 as necessary conditions.

4.2.3 A Set of Independently Sufficient Conditions. If a set of conditions contains alternatives, any one or more of which is deductively sufficient, as a matter of law, to prove the conclusion, but there may be additional sufficient conditions other than those specified, then this is a **weak disjunction** (indicated by placing the symbol “V” before each condition of the set). The legal rule states that if any one or more of the conditions in the set is true, then this is sufficient proof for establishing the truth of the conclusion, as a matter of law (i.e., the conclusion must be true, as a matter of legal authority).

For example, in the statutory sentence quoted in 4.2.1 above, if we were interpreting this sentence in isolation, we would use weak disjunctions to represent the italicized tokens of “or” in the following phrases: “personal injury suffered *or* disease contracted,” “in the active military, naval, *or* air service,” and “discharged *or* released.” More context or interpretation would be required before we would be warranted in using a strong disjunction (see next discussion). Also, in Figure 1, each of the three main conditions (preceded by “&■”) can be proved by evidence of some disability, of which PTSD is only one type. Thus, the weak disjunction indicates an open-ended list of possible disabilities (“V [2 of ...] ...”).

4.2.4 A Necessary Set of Independently Sufficient Conditions. If a set consists of conditions that are independently sufficient, as a matter of law, to deductively prove the conclusion, but the conclusion can be true only if at least one of those conditions is true, then this is a **strong disjunction** (indicated by placing the symbols “■V” before each of the sufficient conditions). The addition of the symbol for necessity (“■”) indicates that the set of conditions as a whole is necessary for proving the conclusion. The legal rule states that, as a matter of law, not only does any of the disjoined conditions provide sufficient proof for the truth of the conclusion (i.e., the conclusion must be true, as a matter of legal authority), but also the conclusion can be true only if at least one of the conditions in the set is true (the truth of the conclusion entails the truth of the disjunction of the conditions).

We may encounter such strong disjunctions when the legal rule states a necessary set of alternative conditions for a valid legal claim, or when a set of alternative conditions occurs in an exception to a broader rule. In addition, judicial decisions interpreting statutes or regulations often focus on only one rule condition, without enumerating the entire set of jointly sufficient conditions, and may announce that it is a necessary condition for proving the conclusion. We find an example in *Rollings v. Brown*, 8 Vet.App. 8 (1995), where the court stated:

For a veteran to be service connected for a disease or injury, the appellant must demonstrate (1) that the disease or injury was incurred in or aggravated by service, 38 U.S.C. § 1110, 38 C.F.R. § 3.303 (1994); or (2) that a certain disease or injury manifested to a degree of 10% or more within one year post-service.

Rollings, 8 Vet.App. at 12. This sentence states that proving service connection entails proving one or both of these conditions.

4.2.5 Merely Relevant Factors and Irrelevant Factors. Some legal rules do not specify a deductive inferential relationship from conditions to conclusion, but only a probabilistic one. In U.S. law, a typical test for determining whether evidence is “**relevant**” is if the evidence “has any tendency to make a fact [of consequence in determining an action] more or less probable than it would be without the evidence” (Federal Rule of Evidence 401). Conditions that are merely relevant to proving a conclusion are often said to be “relevant factors” for drawing the conclusion. We indicate a merely relevant condition by placing the symbol “**RF**” before the condition in the rule tree. A relevant factor or condition, if true, merely influences the probability that a conclusion is true (if the condition is supporting) or false (if the condition is opposing). A legal rule that establishes a condition as relevant establishes the relevancy as a matter of law or legal authority. Authoritatively determining what evidence is relevant for drawing a conclusion authoritatively influences the semantics of legal terminology.

For example, in *Maxson v. Gober*, 230 F.3d 1330 (Fed.Cir. 2000), the court stated:

The court held, and we agree, that evidence of a prolonged period without medical complaint can be considered, along with other factors concerning the veteran's health and medical treatment during and after military service, as evidence of whether a pre-existing condition was aggravated by military service.

Maxson, 230 F.3d at 1333. Under this rule, the fact-finder is entitled, as a matter of law, to take into account “evidence of a prolonged period without medical complaint” when determining “whether a pre-existing condition was aggravated by military service.”

Occasionally, a legal rule might specify that a factor or condition is **irrelevant** to drawing a conclusion, as a matter of law (which we indicate by placing “**IRF**” before the condition). This is a command that the condition must not affect the probability of the conclusion, and is generally interpreted by decision makers as a command that such evidence is not to be considered at all in the decision making. Authoritatively determining what evidence is irrelevant for drawing a conclusion also authoritatively influences the semantics of legal terminology.

For example, the following sentence is from *Godwin v. Derwinski*, 1 Vet.App. 419, 427 (1991):

Furthermore, the Board's analysis designed to show that the arthritic condition is not “attributable to any incident of service”, *Godwin*, BVA 90–10507, at 10, is irrelevant under the presumption established by 38 U.S.C. § 312(b)(12) (1988) ...

Here, the court authoritatively holds that the Board's analysis must have no effect on the inference authorized by the presumption.

4.2.6 A Rebutting Condition. The first of two opposing inferential relations is **rebuttal**, in which the truth of the defeating condition is deductively sufficient, as a matter of law, to disprove the conclusion (i.e., to prove that the conclusion is false). [24] We indicate rebuttal by placing “**REBUT**” before the condition. The rule states that if the defeating condition is true, then as a matter of law the conclusion must be false. If the defeating condition is false or undecided, then the conclusion retains whatever truth-condition it would have had in the absence of the defeater. Such a rebuttal relation is often found in what is called an “affirmative defense” to a claim – i.e., a defensive argument that is sufficient, as a matter of law, to defeat the claim.

For example, the sentence from 38 U.S.C.A. § 1110 quoted in 4.2.1 above concludes as follows:

... the United States will pay to any veteran thus disabled ... compensation as provided in this subchapter, but no compensation shall be paid if the disability is a result of the veteran's own willful misconduct or abuse of alcohol or drugs.

The complex clause introduced by “but” creates a defeating condition: “the disability is a result of the veteran's own willful misconduct or abuse of alcohol or drugs.” We know that this is a rebutting condition because its truth explicitly implies the negation of the compensation conclusion (“no compensation shall be paid”).

Another context in which rebuttals can occur is when the legal rule states a legal presumption. The general form of a presumption is: **if** some proposition (the so-called “basic fact”) is proved to be true, **then** some other proposition is presumed to be true (the “presumed fact”), **unless** some third proposition is true that defeats the working of the presumption. [31] When the “unless” is interpreted as rebuttal, then the presumption has a **strong defeater**, because the truth of the defeater is sufficient to deductively disprove the conclusion.

For example, consider the following sentences from a regulation:

Previous determinations which are final and binding, including decisions of service connection, ... will be accepted as correct in the absence of clear and unmistakable error. Where evidence establishes such error, the prior decision will be reversed or amended.

38 C.F.R. § 3.105(a) (2017) (Revision of decisions). More formally, this presumptive rule is: “**if** there is a previous determination that is final or binding, including decisions of service connection, **then** that determination will be accepted as [will be presumed to be] correct, **unless** evidence establishes that the determination involved clear and unmistakable error.” This presumption includes a strong or rebutting defeater because, if the defeater is true, then the previous determination will be reversed as incorrect (i.e., not accepted as correct).

4.2.7 An Undercutting Condition. The second opposing inferential relation is **undercut**, in which the truth of the defeating condition does not rebut or disprove the conclusion itself, but merely undermines the presumptive support for the conclusion. [24] We indicate undermining by placing “**UNDC**” before the condition. The rule states that if the defeating condition is true, then as a matter of law the presumption ceases to operate. If the defeating condition is false or undecided, then the conclusion retains whatever truth-condition it would have had in the absence of the defeater. When the “unless” in a presumption is interpreted as undermining, then the presumption has a **weak defeater**, because the truth of the defeater merely undermines the presumptive method of proving the conclusion.

For example, consider the following sentence from the statute:

... in the case of any veteran who served for ninety days or more during a period of war ... a chronic disease becoming manifest to a degree of 10 percent or more within one year from the date of separation from such service ... shall be considered to have been incurred in or aggravated by such service, notwithstanding there is no record of evidence of such disease during the period of service.

38 U.S.C.A. § 1112(a)(1) (2017). This “presumption of service connection” would be represented as having a weak defeater.

When “there is affirmative evidence to the contrary,” this merely precludes the claimant from relying on the presumption, but leaves open the possibility of directly establishing a service connection through other evidence. *Goodsell v. Brown*, 5 Vet.App. 36, 44-45 (1993); see 38 C.F.R. § 3.303(d) (2017) (Principles relating to service connection). As this example illustrates, it sometimes requires extraction of information from statutes, regulations, and judicial decisions in order to accurately represent the complete system of governing legal rules.

5 MAKING THE FACT-FINDING REASONING COMPUTABLE

A completed rule tree of the kind discussed in the previous section provides an authoritative list of rule conditions, which identify all of the propositions relevant to deciding a disability claim. A BVA decision supplies the findings of fact corresponding to these rule conditions, to determine whether enough of them are satisfied in the specific case so that the veteran claimant is entitled to compensation. This section of the paper begins an analysis of the semantic types needed to extract the fact-finding reasoning from such BVA decisions, so that we can eventually identify successful and unsuccessful patterns of evidentiary argument. For purposes of mining arguments or lines of reasoning without prejudging their form, we employ a minimalist definition of argument. By an “argument” we mean simply sets of propositions, one of which (the conclusion) can be reasonably believed to be true if the other propositions (the premises or conditions) are reasonably believed to be true [11].

Because sentences can express propositions that form arguments, identifying the inferential roles of individual sentences is a critical task in semantic analysis. Identifying those roles within adjudicatory decisions, however, presents special problems. By contrast with other legal documents such as statutes or regulations, adjudicatory decisions have a wide diversity of roles for sentences. For example, such decisions generally state the legal rules applicable to the decision, as well as legal policies and principles, and provide citations to authority. They also state the procedural history of the case, state the rulings on motions made by the attorneys for the parties, explain the bases for those rulings, and announce the final decisions in the case. Generally, they also summarize the evidence presented and the arguments of the parties based on that evidence, and they state and explain the tribunal’s findings of fact based on that evidence. While law students in the United States can become somewhat proficient at identifying sentence roles, this is a skill that is difficult to learn and particularly difficult to automate.

While the sentence roles that we discuss in this section are familiar to attorneys and judges, the annotation of sentences within legal documents for their inferential role is quite difficult. It underscores the need to develop protocols that provide methods and criteria for manually annotating texts, and a set of conventions governing the generation of data. In the LLT Lab, we develop protocols for assigning roles in four stages: (I) collecting the wide variety of examples of normal forms, linguistic variants, and aberrant forms of the semantic type; (II) deriving from those examples some general guidelines or criteria for manually annotating these types within texts; (III) deriving from the general guidelines a flow chart or decision tree for manually making annotations; and (IV) developing programming rules for automating the annotation process as much as possible. This

section of the paper discusses the roles we have found useful when annotating the fact-finding portions of adjudicatory decisions.

5.1 Semantic Types or Roles for Sentences or Clauses in Adjudicatory Decisions

We have developed and found useful the ten semantic types of sentences or clauses listed in Table 2. This section describes each type briefly, and provides examples from appellate decisions (5.1.1 to 5.1.6) and from decisions of fact-finding tribunals (5.1.7 to 5.1.10).

5.1.1 Citation Sentence or Clause. A citation sentence is a sentence whose primary function is to reference legal authorities or other materials, and which usually contains standard notation that encodes useful information about the cited source. For example, a citation sentence might reference a controlling judicial decision, or a statute or regulation, or a transcript that is in the evidentiary record of the case. While “citation” is often a type of sentence, citations are also embedded as clauses in other types of sentences. The information encoded in legal citations often provides important linguistic cues about the types of other sentences – e.g., when a citation to a rule-making institution is embedded in or immediately follows a sentence, the citation provides some evidence of the sentence’s being a legal-rule sentence. An example of a citation sentence is: “*Hansen v. Principi*, 16 Vet.App. 110, 111 (2002) (citing *Caluza v. Brown*, 7 Vet.App. 498, 505 (1995), *affd*, 78 F.3d 604 (Fed.Cir. 1996) (table)).”

Table 2. Ten Semantic Types of Sentence or Clause Found in Judicial Decisions, Based on Reasoning Roles

- | | |
|-------|--|
| i. | Citation sentence or clause |
| ii. | Legal-rule sentence or clause |
| iii. | Legal-policy sentence or clause |
| iv. | Policy-based-reasoning sentence or clause |
| v. | Ruling or holding sentence or clause |
| vi. | Rule-based-reasoning sentence or clause |
| vii. | Evidence sentence or clause |
| viii. | Finding-of-fact sentence or clause |
| ix. | Evidence-based-reasoning sentence or clause |
| x. | Procedural-fact sentence or clause |

5.1.2 Legal-Rule Sentence or Clause. A legal-rule sentence is a sentence that primarily states one or more legal rules in the abstract, without stating whether the conditions of the rule(s) are satisfied in the case being decided. Such sentences can provide important information for constructing a rule true, and provide building blocks for arguments about issues of fact. From the perspective of argument mining, we are interested not only in identifying rules that are binding in the jurisdiction, but also rules of other jurisdictions that might still play a persuasive role.

An example of a sentence stating a legal rule is:

The Court of Appeals for Veterans Claims has correctly noted that in order to establish service connection or service-connected aggravation for a present disability the veteran must show: (1) the existence of a present disability; (2) in-service incurrence or aggravation of a disease or injury; and (3) a causal relationship between the present disability and the disease or injury incurred or aggravated during service.

Shedden v. Principi, 381 F.3d 1163, 1166-67 (Fed. Cir. 2004).

The annotation of legal-rule sentences in BVA decisions illustrates the need for and development of protocols. One predictor for being a legal-rule sentence is either containing a citation clause that references a rule-making authority, or being followed immediately by such a citation sentence. If the citation refers to a rule-making authority that has jurisdiction over the BVA (e.g., the Federal Circuit), then this is further evidence of a binding rule. In addition, if the target sentence contains explicit linguistic cues by which we can attribute an embedded general proposition to that rule-making authority [33], this provides further evidence that it is a legal-rule sentence. When such criteria are systematically assembled, they provide a protocol for annotating legal-rule sentences.

5.1.3 Legal-Policy Sentence or Clause. A legal-policy sentence is a sentence that primarily states one or more legal policies, principles or objectives in the abstract, without reasoning about whether the policy is being applied correctly in the decided case. The legal policies might derive from the legislature in enacting a statute, or from an administrative agency when engaged in rulemaking or adjudication, or from a court when deciding a case. Legal policies provide important foundations for policy-based reasoning that concludes with adopting or rejecting a legal rule. For example, an explicit statement of policy cited in a regulation is: "It is the defined and consistently applied policy of the Department of Veterans Affairs to administer the law under a broad interpretation, consistent, however, with the facts shown in every case." 38 C.F.R. § 3.102. An example of a legal-policy sentence in a judicial decision is: "To the extent that Congress has relaxed evidentiary requirements in the VA context, it did so to benefit, not penalize, claimants." *AZ v. Shinseki*, 731 F.3d 1303, 1322 (Fed. Cir. 2013).

5.1.4 Policy-Based-Reasoning Sentence or Clause. A policy-based-reasoning sentence is a sentence that primarily applies legal policies to decide legal issues – often to decide whether to adopt or reject a legal rule, but occasionally to decide a finding of fact. Where the issue to be decided is which legal rule to adopt or reject, the policy-based reasoning that justifies or explains the decision might exhibit various traditional patterns (e.g., applying canons of statutory construction or the principle of *stare decisis*), or may invoke policy balancing, means-end causal reasoning, or principle-rule consistency. An example of policy-based reasoning is: "Penalizing assault victims for that failure [to report the assault] would hardly comport with a system in which "the importance of systemic fairness and the appearance of fairness carr[y] great weight." *AZ v. Shinseki*, 731 F.3d 1303, 1322 (Fed. Cir. 2013).

5.1.5 Ruling or Holding Sentence or Clause. A sentence that states a ruling or holding is a sentence that primarily states, "as a matter of law," whether some particular legal rule is satisfied in the case, and applies the rule to the specifics of the case. Such a decision might be about either process or substantive issues. Such issues are often raised for decision by a motion or objection, and the court (judge as judge) makes the resulting decision, not the

factfinder. A critical characteristic is that under certain circumstances a ruling or holding can be treated as a legal rule, applicable to other cases involving similar circumstances. An example of a sentence stating a ruling or holding as a matter of law is: "The Court holds that the Secretary's proposed interpretation of the regulation, which mirrors the relevant statute, 38 U.S.C. § 1101, to broadly exclude tinnitus is not persuasive." *Fountain v. McDonald*, (27 Vet. App. 258, 259-60 (2015)). Another example from the same appellate decision is: "In making this determination, the Board failed to provide an adequate statement of reasons or bases in light of the governing caselaw." *Id.* at 274.

5.1.6 Rule-Based-Reasoning Sentence or Clause. A sentence that states rule-based reasoning is a sentence that reasons from a foundation of legal rules and facts to a ruling or holding as a matter of law (i.e., a conclusion of law). Such sentences often report deductive reasoning (such as applying *modus ponens* or *modus tollens* to the conditional structure of a legal rule), but they might also report probabilistic or defeasible reasoning. Rule-based reasoning can also occur within a passage of evidence-based reasoning, as when the fact-finder applies a presumptive rule of law in assessing the probative value of the evidence. An example of rule-based reasoning leading to a conclusion of law is: "In his brief, however, Mr. Buie makes no arguments with respect to his claim for benefits for post-traumatic stress disorder, and the Court therefore deems any appeal with respect to that claim abandoned." *Buie v. Shinseki*, 24 Vet. App. 242, 243 (2010). This sentence is immediately followed by this citation sentence: "*See Grivois v. Brown*, 6 Vet.App. 136, 138 (1994) (holding that issues or claims not argued on appeal are considered abandoned)."

5.1.7 Evidence Sentence or Clause. An evidence sentence is a sentence that primarily states the content of the testimony of a witness, states the content of documents introduced into evidence, or describes other evidence. Evidence sentences provide part of the foundation either for findings of fact or for rulings of law. An example of a statement of evidence is: "The clinician opined that 'It is LESS LIKELY THAN NOT [sic] that the Veteran's claimed mental health condition, to include PTSD, incurred, was caused by or proximately due to his 6 months in military service.'" Board of Veterans Appeals #1340434 (12/06/2013). Although this sentence was written by the Board, it attributes a proposition to a clinician, and reports evidence rather than the Board's findings.

5.1.8 Finding-of-Fact Sentence or Clause. A finding-of-fact sentence (an evidence-based-finding sentence, or simply finding sentence) is a sentence that primarily states an authoritative finding, conclusion or determination of the trier of fact – a decision made "as a matter of fact" instead of "as a matter of law." If the finding is by a judge or administrative law judge, it is normally called a "finding," whereas a factual determination by a jury is normally called a "verdict." An example of a statement of a finding of fact is: "The most probative evidence fails to link the Veteran's claimed acquired psychiatric disorder, including PTSD, to active service or to his service-connected residuals of frostbite." Board of Veterans Appeals #1340434 (12/06/2013). This finding of fact example is founded in part on the evidence example in 5.1.7.

5.1.9 Evidence-Based-Reasoning Sentence or Clause. An evidence-based-reasoning sentence is a sentence that primarily reports the trier of fact's reasoning in making the findings of fact. Such reasoning often involves an assessment of the credibility and probative value of the evidence, and may also include application

of substantive or process rules, and occasionally even legal policies. An example of a statement of evidence-based reasoning is: “Also, the clinician’s etiological opinions are credible based on their internal consistency and her duty to provide truthful opinions.” Board of Veterans Appeals #1340434 (12/06/2013). This states part of the reasoning leading from the evidence example in 5.1.7 to the finding of fact example in 5.1.8.

5.1.10 Procedural-Fact Sentence or Clause. A procedural-fact sentence is a sentence that primarily states one or more procedural facts about the specific case, such as what motions were filed or the disposition of the case at the trial level. Such sentences might therefore provide important foundational information for the rulings or holdings that are made in the specific case. Very often the court takes official notice of such facts, and they are not disputed by the parties. An example of a statement of procedural fact is: “This matter comes before the Board of Veterans’ Appeals (Board) on appeal from a July 2008 rating decision of the Department of Veterans Affairs (VA) Regional Office (RO) in Montgomery, Alabama.” Board of Veterans Appeals #1340434 (12/06/2013).

5.2 Frequencies of Sentence Types in a Sample of BVA PTSD Disability Decisions

We find all ten types of sentence in adjudicatory decisions, whether those decisions are issued by courts or by administrative tribunals, and whether by appellate or trial courts or tribunals. Even in the case of a tribunal that possesses only fact-finding authority (i.e., has no authority to adopt legal rules that are binding in other cases or on other tribunals), that tribunal’s

decisions generally recite the governing rules and policies. On the other hand, reviewing courts with no fact-finding authority often issue decisions that contain sentences that state the admitted evidence, the evidentiary reasoning of the lower tribunal, and the findings of fact made in that lower tribunal. The scope of adjudicatory authority, therefore, does not automatically indicate the types of sentence occurring in the decisions.

Table 3 shows frequencies for certain types of sentence in a sample of 20 BVA decisions. This set of decisions was chosen arbitrarily, as an initial dataset for developing our annotation protocols. Table 3 illustrates the wide range of number of sentences per decision (from a low of 54 to a high of 921). About 18% of those sentences were citations, and nearly 13% were legal-rule sentences. While the number of citation sentences in each decision was greater than the number of legal-rule sentences, in some decisions the difference was quite small. We also identified the key findings of fact directly related to the three PTSD rule conditions in Figure 1, and determined that they comprised only about 4% of the sentences. Many legal-rule sentences that state the conditions for proving the service-connectedness of PTSD use the same terminology as sentences that state the relevant evidence, the Board’s reasoning, and the Board’s findings. Thus, it is difficult to distinguish among those different types of sentence. The difficulty of correctly and consistently identifying such finding sentences out of the total number of sentences, and of distinguishing them from legal-rule and other types of sentence, underscores the importance of protocols.

Table 3. Frequencies of Sentence Types in a Sample of BVA PTSD Disability Decisions

BVA Citation Number	PTSD Claim Outcome	Freq. of Sentences	Freq. of Citation Sentences	Freq. of Legal-Rule Sentences	Freq. of PTSD Finding Sentences
1302554	Denied	563	35	32	15
1303141	Denied	287	68	66	14
1315144	Denied	196	38	23	8
1316146	Denied	249	56	47	14
1316336	Denied	216	38	20	14
1334312	Denied	340	60	38	10
1343153	Denied	240	27	18	22
1400029	Denied	616	176	122	11
1413417	Denied	921	96	58	9
1431031	Denied	151	35	14	19
1445540	Remanded	145	30	23	6
1455333	Granted	125	35	21	9
1505726	Denied	425	98	71	17
1514581	Denied	327	39	26	24
1525217	Granted	159	40	32	15
1526599	Denied	153	39	30	8
1608262	Denied	158	28	23	6
1613894	Denied	218	51	28	7
1630016	Granted	131	34	31	7
1630402	Granted	54	10	7	5
Totals		5,674	1,033	730	240

7 CONCLUSION AND FUTURE WORK

The empirical development of the type systems for propositional connectives and for sentence roles, reported above, is an important benchmark in making legal reasoning computational. Through our current and future work, we hope to stimulate a shared research effort to create diverse and publicly available datasets in law, to empirically evolve optimal sets of semantic types for argument mining, and to refine protocols for accurately applying those types to texts. We plan to make the dataset of decisions for veterans' disability claims publicly available for non-commercial research, increase the number of annotated BVA decisions currently available, and add annotated appellate decisions to the dataset. We will continue to refine the current protocols for semantic types, and to test machine learning and rule-based programming for automating the annotation task. We are investigating our working hypothesis that rule trees and sentence typing lead to algorithms for searching for, retrieving and re-ranking decisions from the dataset in response to queries about legal reasoning. Finally, we are working to develop a semantic type system for patterns of evidentiary reasoning [34], so we can monitor successful and unsuccessful patterns of argument, as a means of making claims processes more transparent, efficient and accurate.

ACKNOWLEDGMENTS

The authors wish to thank the Maurice A. Deane School of Law at Hofstra University for its support of the Research Laboratory for Law, Logic and Technology and of this project. Special thanks to Patricia A. Kasting for her research assistance on this project.

REFERENCES

- [1] Administrative Office of the United States Courts. 2015. Federal Judicial Caseload Statistics 2015.
- [2] Al-Abdulkarim, L., Atkinson, K., and Bench-Capon, T. 2016. Statement Types in Legal Argument. In *Legal Knowledge and Information Systems (JURIX 2016)*, Bex, F., and Villata, S., eds. IOS Press, 3-12.
- [3] Allen, M.P. 2007. Significant Developments in Veterans Law (2004-2006) and What They Reveal about the U.S. Court of Appeals for Veterans Claims and the U.S. Court of Appeals for the Federal Circuit. *University of Michigan Journal of Law Reform* 40: 483-568. University of Michigan.
- [4] Ashley, K.D., and Brüninghaus, S. 2009. Automatically classifying case texts and predicting outcomes. *Artificial Intelligence and Law* 17: 125-165.
- [5] Ashley, K.D., and Walker, V.R. 2013a. Toward Constructing Evidence-Based Legal Arguments Using Legal Decision Documents and Machine Learning. In *Proceedings of the 14th International Conference on Artificial Intelligence and Law (ICAIL 2013)*. ACM, New York, NY, 176-180.
- [6] Ashley, K.D., and Walker, V.R. 2013b. From Information Retrieval (IR) to Argument Retrieval (AR) for Legal Cases: Report on a Baseline Study. In *Legal Knowledge and Information Systems, Ashley, K. D., Ed.* IOS Press, 29-38.
- [7] Bansal, A., Bu, Z., Mishra, B., Wang, S., Ashley, K., and Grabmair, M. 2016. Document Ranking with Citation Information and Oversampling Sentence Classification in the LUIIMA Framework. In *Legal Knowledge and Information Systems (JURIX 2016)*, Bex, F., and Villata, S., eds. IOS Press, 33-42.
- [8] Board of Veterans' Appeals, U.S. Department of Veterans Affairs. 2015. Annual Report, Fiscal Year 2015.
- [9] Dragoni, M., Governatori, G., and Villata, S. 2015. Towards Automated Rules Generation from Natural Language Legal Texts. *Workshop on Automated Detection, Extraction and Analysis of Semantic Information in Legal Texts (ASAIL-2015)*.
- [10] Francesconi, E. 2009. An Approach to Legal Rules Modelling and Automatic Learning. In *Legal Knowledge and Information Systems (JURIX 2009)*, Governatori, G, ed. IOS Press, 59-68.
- [11] Francesconi, E., Montemagni, S., Peters, W., and Tiscornia, D. 2010. Integrating a Bottom-Up and Top-Down Methodology for Building Semantic Resources for the Multilingual Legal Domain. In *Semantic Processing of Legal Texts, LNAI 6036*, 95-121. Springer: Berlin.
- [12] Grabmair, M., Ashley, K.D., Chen, R., Sureshkumar, P., Wang, C., Nyberg, E., and Walker, V.R. 2015. Introducing LUIIMA: An Experiment in Legal Conceptual Retrieval of Vaccine Injury Decisions Using a UIMA Type System and Tools. In *Proceedings of the 15th International Conference on Artificial Intelligence & Law (ICAIL 2015)* 69-78. ACM, New York.
- [13] Grabmair, M., Ashley, K.D., Hwa, R. and Sweeney, P.M. 2011. Toward Extracting Information from Public Health Statutes using Text Classification and Machine Learning. In *Legal Knowledge and Information Systems (JURIX 2011)*, Atkinson, K., ed. IOS Press, 73-82.
- [14] Hachey, B., and Grover, C. 2006. Extractive summarization of legal texts. *Artificial Intelligence and Law* 14:305-345.
- [15] Hage, J., and Verheij, B. 1994. Reason-Based Logic: A Logic for Reasoning with Rules and Reasons. *Law, Computers & Artificial Intelligence* 3(2-3): 171-209.
- [16] Jackson, P., Al-Kofahi, K., Tyrrell, A., and Vachher, A. 2003. Information extraction from case law and retrieval of prior cases. *Artificial Intelligence* 150:239-290.
- [17] Maat, E. de, Krabben, K., and Winkels, R. 2010. Machine Learning versus Knowledge Based Classification of Legal Texts. In *Legal Knowledge and Information Systems (JURIX 2010)*, Winkels, R.G.F., ed. IOS Press, 87-96.
- [18] Markovich, R., Syi, and Hamp, G. 2015. Elliptical Lists in Legislative Texts. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law (ICAIL '15)* 192-195. ACM, New York.
- [19] McCarty, L.T. 2007. Deep Semantic Interpretations of Legal Texts. In *Proceedings of the 11th International Conference on Artificial Intelligence and Law (ICAIL '07)* 217-224. ACM, New York.
- [20] Mochales, R., and Moens, M-F. 2011. Argumentation mining. *Artificial Intelligence and Law* 19: 1-22. Springer.
- [21] Moens, M-F., Boiy, E., Mochales, R., and Reed C. 2007. Automatic Detection of Arguments in Legal Texts. In *Proceedings of the 11th International Conference on Artificial Intelligence and Law (ICAIL '07)* 225-230. ACM, New York.
- [22] Moshiahswili, V.H. 2015. The Downfall of Auer Deference: Veterans Law at the Federal Circuit in 2014. *American University Law Review* 64: 1007-1087. American University.
- [23] Opsomer, R., De Meyer, G., Cornelis, C., and Van Eetvelde, G. 2009. Exploiting Properties of Legislative Texts to Improve Classification Accuracy. In *Legal Knowledge and Information Systems (JURIX 2009)*, Governatori, G, ed. IOS Press, 136-145.
- [24] Pollock, J.L. 1990. *Nomic Probability and the Foundations of Induction*. Oxford University Press, New York.
- [25] Prakken, H., and Sartor, G. 1996. A dialectical model of assessing conflicting arguments in legal reasoning. *Artificial Intelligence and Law* 4: 331-368.
- [26] Saravanan, M, and Ravindran, B. 2010. Identification of rhetorical roles for segmentation and summarization of a legal judgment. *Artificial Intelligence and Law* 18:45-76.
- [27] Savelka, J., and Ashley, K.D. 2015. Transfer of Predictive Models for Classification of Statutory Texts in Multi-jurisdictional Settings. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law (ICAIL 2015)* 216-220. ACM, New York.
- [28] Sergot, M.J., Sadri, F., Kowalski, R.A., Kriwaczek, F., Hammond, P., and Cory, H.T. 1986. The British Nationality Act as a Logic Program. *Communications of the ACM* 29(5): 370-386.
- [29] U.S. Court of Appeals for Veterans Claims. 2015. Annual Report, October 1, 2014 to September 30, 2015 (Fiscal Year 2015).
- [30] Verheij, B., Hage, J., and Lodder, A.R. 1997. Logical Tools for Legal Argument: A Practical Assessment in the Domain of Tort. In *Proceedings of the Sixth International Conference on Artificial Intelligence and Law (ICAIL 1997)*. ACM, New York, 243-249.
- [31] Walker, V.R. 2014. Representing the use of rule-based presumptions in legal decision documents. *Law, Probability and Risk* 13(3-4): 259-275.
- [32] Walker, V.R., Carie, N., DeWitt, C.C., and Lesh, E. 2011. A framework for the extraction and modeling of fact-finding reasoning from legal decisions: lessons from the Vaccine/Injury Project Corpus. *Artificial Intelligence and Law* 19: 291-331.
- [33] Walker, V.R., Bagheri, P., and Lauria, A.J. 2015. Argumentation Mining from Judicial Decisions: The Attribution Problem and the Need for Legal Discourse Models. *Workshop on Automated Detection, Extraction and Analysis of Semantic Information in Legal Texts (ASAIL-2015)*.
- [34] Walker, V.R., Vazirova, K. and Sanford, C. 2014. Annotating Patterns of Reasoning about Medical Theories of Causation in Vaccine Cases: Toward a Type System for Arguments. In *Proceedings of the First Workshop on Argumentation Mining (Baltimore, Maryland, June 26, 2014)* 1-10. ACL.
- [35] Winkels, R., and Hoekstra, R. 2012. Automatic Extraction of Legal Concepts and Definitions. In *Legal Knowledge and Information Systems (JURIX 2012)*, Schäfer, B., ed. IOS Press, 157-166.
- [36] Wyner, A., and Peters, W. 2011. On Rule Extraction from Regulations. In *Legal Knowledge and Information Systems (JURIX 2011)*, Atkinson, K.M., ed. IOS Press, 113-122.