NORTHWESTERN UNIVERSITY

Extracting and Applying Legal Rules from Precedent Cases

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

For the degree

DOCTOR OF PHILOSOPHY

Field of Computer Science

By

Joseph Alan Blass

EVANSTON, ILLINOIS

March 2023

ABSTRACT

Extracting and Applying Legal Rules from Precedent Cases

Joseph Alan Blass

In the Western Common Law tradition, legal decisions constrain and guide future cases involving the same legal issues. Although the legal academy disagrees about the specific nature of such precedential reasoning, including about the role of analogy in legal reasoning, there is ample evidence from cognitive science that analogical learning and reasoning are ubiquitous in human cognition. This thesis presents a computational approach to performing legal precedential reasoning and argumentation using analogical learning and reasoning, grounded in research in Artificial Intelligence, Cognitive Science, and Law.

The thesis presents a dataset of historical Illinois intentional tort cases upon which the rest of the thesis is trained and tested, and an algorithm for supervised analogical learning that is useful in instances where it is hard to discern whether an analogical match is useful. It demonstrates that this algorithm can learn, from across a body of case law, schemas that capture the legal information governing those cases. The thesis then presents three algorithms for legal reasoning and prediction using such learned legal schemas: one that reasons directly by analogy from prior cases and legal schemas to a new case, one that reasons *about* the analogies drawn from legal schemas to a new case, and one that converts those schemas into logical rules and reasons about the new case using logic. It also presents a legal argumentation system adapted from the rule-reasoning system. The thesis demonstrates that abstract legal information can be captured through a process of analogical learning, that analogical reasoning can be used to resolve common law cases, and that schemas induced through an analogical learning process can be converted into rules useful for rule-based legal reasoning.

ACKNOWLEDGEMENTS

There are innumerable people who supported and contributed to my work and life across my graduate studies, but here are some numerable *sine qua non*s.

My advisor, Ken Forbus, took a chance on me when I first applied as a Masters' student relatively new to computer science, and again when he offered me a spot in the PhD program. He has supported and encouraged me across these many years as I charted an unusual path that took on a new form more than twice. I have benefited tremendously from his depth of knowledge and guidance, his short- and long-term patience, and his tolerance for my pursuing my passion even at the expense of a speedy graduation. I have also been incredibly fortunate to have the best and most supportive committee—in addition to Ken—that I could have asked for: Dedre Gentner, Sarah Lawsky, and Bryan Pardo. Each of them gave generously of their time, wisdom, and perspective, sharing not only their substantive knowledge of their fields, but also their insights into navigating graduate school, law school, academia generally, and life.

I owe a special debt of gratitude to Shari Diamond, who encouraged me to apply to the JD-PhD program, and without whom I would not have gone to law school. Shari has nurtured and encouraged my development as a budding legal academic, and I know with certainty that I would not be where I am and heading where I am headed without her. A special thanks also to Scott Cambo, who among many other wonderful moments as a friend and fellow student, was the person who made me aware that Northwestern even had a JD-PhD program for me to be interested in.

A big thank you to the many Northwestern professors in addition to those already listed, who took the time to talk to me outside of class—often not about the coursework—and to read my work, including Deborah Tuerkheimer, Andrew Koppelman, Ronald Allen, Janet Brown, Christos

Dimoulas, Doug Downey, Paul Gowder, Tom Hinrichs, Ian Horswill, Jonathon Koehler, James Lindgren, Dan Linna, Doug Medin, James McMasters, Martin Redish, Chris Riesbeck, and James Speta. And thank you to those outside of Northwestern who have similarly been generous with their time and have encouraged my personal and professional development: A severely truncated list includes David Aha, Kevin Ashley, Victoria Nourse, Mathias Scheutz, Josh Tenenbaum, and Rebecca Wexler, but there are too many for me to list.

Thank you to my many friends and peers that made this experience the joy that it has been. I apologize in advance, because after this much time across multiple home bases I am quite certain I have left important people off this list, which I have anyways tried to restrict to those who directly contributed to my graduate educational experience and the development of my ideas: Kat Albrecht, Mike Bancroft, Chandra Bhagavatula, Conor Bidner, Scott Cambo, Maria Chang, Wai Lee Chin Feman, Hillary Chutter-Ames, Janene Cielto, Max Crouse, Noah Cruickshank, Seth Dickinson, Tesca Fitzgerald, Becca Goldstein, Will Hancock, Bryan Head, Kerri Howard, Charlie Isaacs, Subu Kandaswamy, Anne Kostroski, Mukund Kuthalam, Chen Liang, Michael Lucas, Debbie Ma, CJ McFate, Matt McLure, Constantine Nakos, Thanapon Noraset, Carrie Ost, Annie Prossnitz, Irina Rabkina, Danilo Ribeiro, Hecate Robison, Tommy Rousse, Prem Seetharaman, Wil Stevens, Dustin Stonecipher, Ari Tolman, Nick Trimis, Madeline Usher, Matea Varvodic, and many, many more.

Thank you to my family for their love, support, and occasionally aggressive questioning: Bridget, Carlos, Patrick, Ewa, Josiah, Simon, and Evelyn. Words are drastically insufficient to thank my mother and father for all they have done for me. In addition to being fantastic parents, they have brought their tremendous intellects, knowledge, and experience to bear in helping discuss and develop my work over the years.

My greatest gratitude remains for my spouse Bianca, who moved to Chicago with me for what was supposed to be a one-year Masters' program a decade ago. So much has happened since then, and I know I would not have been able to do everything that I did without her encouragement, love, interest, and engagement. I hope she will forgive me for not taking her advice to include a picture of our cats, Herschel and Mika, here in the acknowledgments section, although it's true that they deserve acknowledgment (will this do?). Thank you, Bianca, for being a true partner at every step of this journey.

This research was supported by the Computational Cognition and Machine Intelligence Program of the Air Force Office of Scientific Research under award #FA9550-20-1-0091. This thesis was researched and written on a Microsoft Surface Pro 7, with a 1.3 GHz Intel i7 processor and 16 GB of RAM.

TABLE OF CONTENTS

TABLE OF FIGURES

TABLE OF TABLES

INTRODUCTION

The law can be complicated, and hiring a lawyer can be expensive. Those two facts combine to keep people out of the legal system who might otherwise benefit from having access to it. Because the law is complicated, people who think they may have been wronged and want to know if they have legal recourse to redress it often have to hire a lawyer just to understand the law's perspective on what happened to them and whether they have a case. Because hiring a lawyer can be expensive, potential litigants without the means to pay for even an initial consultation might be unable to acquire the knowledge they need to understand their rights and obligations. Automating the roles of judges and lawyers is questionable to the point of undesirable, but an automated system that could help analyze the law, explain how it is likely to be applied in a case, and suggest some arguments that might be useful in arguing that case—a system combining elements of the role of a paralegal and lawyer—could help people who would otherwise keep themselves out of the legal system decide whether it is worth investing in their case.

In the Anglo-American Common Law system, legal cases are decided in concordance with their precedents. Unless a judge decides (and is empowered) to change some rule of decision, or a case presents an issue of first impression, the judge should apply the rule that governed prior similar cases to decide the case at bar. Lawyers and litigants understand this and so craft their arguments within the boundaries of those precedents: they argue that the facts of the case are such that the prior rule applied to those facts will lead to their preferred outcome, or where ambiguity in the prior cases allows, they argue about what the actual rule governing the prior case was. Discerning the nature of the rule applied in the prior case, and applying that rule to the case at bar, are thus essential components of both judges' and lawyers' work in a common law system, and are in fact referred to simply as "common law reasoning."

Researchers studying Artificial Intelligence (AI) and Law have built computational models of common law reasoning, using a variety of methodologies and for a variety of purposes. This thesis adds to that scholarly conversation by introducing a new computational common law reasoning system, Precedential Analogical Reasoning and Argumentation for the Legal domain (PARALegal, or just PARA). This system learns common-law legal principles by comparing precedential cases to each other and building schemas from those comparisons. These schemas can be applied to unseen cases by analogy, or converted into logical rules and applied to those unseen cases in formal logic. In so doing, the system is reasoning from the perspective of a judge or a judicial clerk (or perhaps a lawyer in the early stages of litigation) about what legal outcome a case's facts suggest should result from the case. The system can also use those rules to argue the case from a lawyer's perspective, arguing facts that would lead to the outcome preferred by one side or another.

My research builds on prior work in AI, in Law, and in Cognitive Science. It is the first automated legal reasoning system to model legal case-based reasoning using the Structure Mapping Engine model of analogical reasoning (SME (Forbus, Ferguson, Lovett, & Gentner, 2017), based on Gentner's 1983 Structure Mapping Theory), or to learn legal principles by creating schemas through analogical case-comparison (using SAGE, the Sequential Analogical Generalization Engine (Kandaswamy & Forbus, 2012)). While this thesis is not cognitive modeling research, lawyers and judges are humans and legal reasoning and argumentation is generally the domain of human cognition, so taking inspiration from human cognition and relying on tools built to replicate principles of human cognition makes sense. A computational model of human analogical reasoning seems an especially good fit for common law reasoning, because analogies are a hallmark of common-law argument and written analysis (although, as Research Background demonstrates, the role of analogy in common law reasoning is controversial).

The system incorporates two models of precedential legal reasoning. In one model, rules are extracted from precedent cases, and those extracted rules are applied to the case at bar. In the other, precedent cases (perhaps generalized/schematized precedent cases) are applied to the case at bar, and the rule is only revealed, if at all, after applying the *reasoning* (not the *rule*) of the precedent case to the case at bar. This thesis demonstrates that these approaches can perform better than off-the-shelf machine learning approaches, while producing decisions that are inspectable and explainable by a human user. Furthermore, the same system is used to model what the precedents in the case base indicate the outcome of the case *should* be, as well as how the lawyers for each side might *argue* for their preferred outcome. That is, the same basic system can be used to predict outcomes in cases, as judges do, and to generate arguments for those cases, as lawyers do.

Though it is neither the first system to use explicit symbolic representations of cases, nor the first to start from cases' English-language statements of facts, PARALegal is the first system to marry these two approaches, using a semantic interpreter (the Companions Natural Language Understanding system, or CNLU, (Tomai & Forbus, 2009)to extract predicate logic representations from cases without hand-generating any representations of case facts. Relatedly, while most prior approaches present domain-general theories, the models have been built in domain-specific ways. That is, the models have always involved human-engineered domain-specific legal representations developed for the particular legal domain being considered (e.g., trade secrets, workers' compensation, first-capture rules of property rights). The present research only involved testing PARALegal on the domain of Torts, but it involved no hand-encoding of domain-specific legal concepts or rules. The cases were represented simply as their facts, with no specialized legal knowledge added either directly to case facts or as an intermediate reasoning step. So while it has only been tested on Tort cases, the system should without modification be able to be brought to bear on other domains; it awaits only the dataset upon which to do so.

A. RESEARCH CONTRIBUTIONS

This thesis advances the fields of AI, Law, and the interdisciplinary field of AI & Law in the following ways. First, it introduces a supervised analogical learning algorithm, which allows researchers to learn relational schemas from datasets of cases that lack the deep structural representations that analogical reasoning generally depends on. While this thesis only demonstrates the advantages of this new learning algorithm on the legal domain, this is a pure AI contribution that should be useful in learning from any such dataset, and may prove to be useful even in learning analogical schemas from datasets that have those deep structural representations.

The thesis also presents a new framework for thinking about the role of analogy in learning legal content from precedent cases. While members of the legal academy have debated the role of analogies in resolving some new case by analogy to a prior case, there has been little attention paid to how analogical learning may be the mechanism by which legal rules are learned from a body of case law. This thesis presents a functional implementation of this newly proposed framework, learning symbolic representations of legal rules directly from the facts of the case, without relying on humans hand-encoding the cases or the rules.

The thesis introduces three algorithms to perform legal precedential reasoning, which contribute separately and together to advancing research in AI & Law. The first algorithm implements a purely analogical model of legal precedential reasoning, and is the first algorithm to use Structure Mapping as the basis for legal reasoning and decision-making. The second algorithm reasons about a new case by analogy to schemas of legal doctrines, using Structure Mapping rather than unification to determine whether elements of a claim are present in a new case, which allows for more flexibility in reasoning than does logical deduction. These first two algorithms together describe an alternate vision for the role and nature of analogy in legal reasoning. The third algorithm converts legal analogical schemas into logical rules and reasons with them deductively. Taken together, these three algorithms describe a cycle by which legal rules are developed and discerned across the process of extracting and applying legal knowledge from a body of precedent cases.

The thesis examines the risks presented by legal automation and develops an argument against legal automation. It also examines what qualities an AI ought to have before it is permitted to perform any sort of legal reasoning that might have a real-world impact. Taken together, these provide the tools to determine whether some automated reasoner should be rejected out of hand for some legal reasoning task because it lacks the threshold requirements for performing it.

Finally, the thesis makes publicly available a dataset of historical legal Tort cases encoded in predicate logic. While there are several large, publicly available datasets of judicial cases that have been annotated with metadata about, e.g., what claim was at issue and who won, this is the first publicly available dataset of such cases represented entirely in predicate logic, which will be useful to any AI & Law researchers looking to train or test symbolic legal reasoning systems.

B. OVERVIEW OF THE THESIS DOCUMENT

The rest of this document is organized as follows.

Chapter 2: Research Background. This chapter describes the research background of this thesis. This includes research and theory in Artificial Intelligence, Law, the intersection of both (AI & Law), and Psychology.

Chapter 3: How To Approach Modeling Legal Reasoning. This chapter explains in more depth the goals of the system, and how those goals informed why certain engineering approaches were taken and others rejected. It proposes a new—to the legal academy, at least—role for analogy in legal reasoning, namely, as the mechanism by prior cases are compared *to each other* in order to discern the rules governing those cases. It discusses potential roles for legal automated reasoning systems, what properties those systems should have (particularly transparency and explainability), the utility of cognitive modeling, and the desired scalability of the PARALegal system.

Chapter 4: The Illinois Intentional Tort Qualitative Dataset. This research project required a dataset of cases that involve purely common-law doctrine (to avoid, during an initial research project, having to perform statutory interpretation) and which are represented in propositional logic. Unable to find such a dataset, I created one: the Illinois Intentional Tort Qualitative Dataset (“the Dataset”). This chapter will describe the Dataset, as well as areas in which it could be expanded.

Chapter 5: Conclusion-verified Analogical Schema Induction. For reasons that will be discussed in this chapter, the cases in the Dataset appear sufficiently dissimilar to each other according to the Structure Mapping Engine’s scoring systems that relying on using SME’s similarity score to control schema assimilation in SAGE was ineffective. I developed a new technique, CASI, to determine when to assimilate cases into a generalization. This chapter describes the technique and situations in which it should prove useful. It also describes the generalizations built using CASI over cases in the Dataset, generalizations used by the rest of the experimental systems described in the thesis.

Chapter 6: Legal Precedential Reasoning with Analogical Generalization, Analogical Reasoning, and Rule Learning. This chapter concerns the three legal precedential reasoning systems I constructed. Two are analogical reasoning techniques used to bring the generalizations learned by CASI (described in the prior chapter) to bear on held-out cases in the dataset. The third is a system to turn the learned analogical generalizations of legal rules into actual logical rules, i.e., Horn clauses. The chapter will describe the algorithms and an evaluation thereof.

Chapter 7: Legal Argumentation. This chapter describes a system that uses the rules generated by the system (described in the previous chapter) for legal argumentation, i.e., a lawyer's argument for an outcome favoring a particular side. It reports the results of a pilot experiment and describes future areas of inquiry.

Chapter 8: Limitations and Future Work. This chapter discusses limitations inherent to using computational legal systems. It discusses limitations related to requiring specified case facts, and to the challenges posed by open-textured terms.

Chapter 9 concludes.

CHAPTER 2: RESEARCH BACKGROUND

This work, though first and foremost an exercise in AI research executed using the tools of a computer scientist, is informed by legal scholarship on precedential and analogical reasoning in the law, and by psychology and cognitive science research on everyday reasoning by analogy.

A. LEGAL REASONING AND CASE-BASED REASONING

This research is conducted against the background of a significant body of legal scholarship concerning the mechanisms of precedential reasoning (i.e., reasoning about a current legal case using a body of previously decided cases), and specifically analogical legal reasoning. Scholars are in wide agreement that judges and lawyers often use analogies in the arguments, opinions, and explanations they make, but these scholars disagree about the role of such analogies in the dispositions of actual cases.[1]

At one extreme are scholars who believe that analogies are used only expressively and not rationally, that is, that the analogies used by legal actors do not help to dispose of the substance of cases, but only to persuade or perhaps explain. Thus, Judge Richard Posner, an influential legal scholar, says that "[r]easoning by analogy as a mode of judicial expression is a surface phenomenon. It belongs not to legal thought, but to legal rhetoric[, and] confuse[s] how judges think with how they talk." (Posner, 2006). Posner argues instead that legal decisions are either settled by clearly established rules or, where the rules do not provide an answer, by policy analysis. Alongside Posner in the camp of legal scholars who believe that reasoning by analogy plays no real role in resolving cases are Larry Alexander and Emily Sherwin. Alexander and Sherwin argue that there are only two real forms of legal reasoning: the "natural model" of reasoning by reference to moral and ethical rules, and the "rule model" of reasoning to established legal rules (Alexander & Sherwin, 2008). "Past results cannot determine the outcomes of new disputes," they write. "Analogical reasoning, as such, is not possible." Part of their criticism of analogy stems from their definitional assumptions: first, they define reasoning as "conscious, language-based deliberation about reasons for the choice ultimately made." Because this definition excludes the semi-conscious, automatic discovery of similarity across cases that is one of the features of reasoning by analogy, it defines away some forms of analogical reasoning as not being "reasoning," without

[1] Note that the discussion of legal analogies in this thesis concerns the role of analogy in legal reasoning as a descriptive matter, not a normative matter: this section discusses scholarly disagreement about whether analogies *actually* resolve cases, not whether they *should* resolve cases. Whether it is jurisprudentially proper for analogy to play any such role is well out of scope of my work.

engaging with the work such reasoning might perform in examining cases according to their precedents. Second, they define analogy as drawing the outcome of a new case directly from the outcome of the past case based only on similarity and without regard to the reasons underlying the previous decision (because they argue that, if there is reliance on the reasons underlying the previous case, the new case is not being decided by analogy).

Further along the spectrum are those scholars who believe jurists engage in analogical reasoning to discern rules from precedent cases, but who do not necessarily think individual cases are resolved by analogy. The first is Edward Levi, who wrote a seminal modern text on legal reasoning (Levi, 1949). Though he does not elaborate on how reasoning by analogy in the law proceeds, he describes it as a three-step process whereby cases are aligned; a rule inherent in the first is discerned; then the rule is applied to the latter case. Similarly, Cass Sunstein has argued that even though any given case is governed by some governing idea (like a rule), "the governing idea is not given in advance and applied to the new case. Instead, analogical reasoning helps identify the governing rule and is indispensable to the identification" (Sunstein, 1996). That is no accident: Sunstein argues that courts intentionally leave judgments and legal practices incompletely theorized, avoiding deciding more than is necessary in a given case and allowing the law to develop over time. Sunstein notes that courts may come to find a past case's outcome depended on factors other than those flagged by the precedent court, and that analogy can play an important role in determining what, in fact, was important to a case's disposition (i.e., what the law is) (Sunstein, 1996). By engaging in the process of analogical reasoning, the principles governing the body of case law can be revealed (Sunstein, 1993).

Finally, Scott Brewer similarly agrees with Levi and Sunstein that analogy plays a role in determining what rules are to be applied to a case, but goes further into depth than either as to how those rules are discerned from the precedent cases (Brewer, 1996). Brewer characterizes the process as one of analogical abduction, where a rule is abduced (i.e., hypothesized and found to be possible on the basis of features in the cases) from precedent cases. Brewer sees the extracted rule as being one that supports the *use* of an analogy from one case to another, rather than the rule *directly* governing the cases. That is, the extracted rule or rules justifies the reliance on an analogy to one precedent case over another, rather than itself disposing of the case through application of a rule. Whether in practice such a method may turn out to functionally be the same as a rule directly disposing of a case is unclear.

Finally, at the other extreme from Posner and Alexander & Sherwin are Lloyd Weinreb and Frederick Schauer & Barbara Spellman. Weinreb argues that cases themselves may be resolved directly by analogy to previous precedent cases (Weinreb, 2016). He grounds this argument in the fact that, as a descriptive matter, lawyers and judges appear to do so all the time. He also argues that any rule that could completely dispose of a case is not, in fact, a rule (since it will have to be specific to that case and no other), and that therefore any rule that *mostly* disposes of a case cannot *entirely* dispose of it because there are particulars to the case not covered by the rule. He therefore argues that, while legal rules naturally exist and apply, they leave a gap between what they resolve about the case and what must be resolved for the case to be disposed of, and analogies directly to precedent cases are necessary to fill that gap and dispose of the case.[2] Schauer and Spellman agree that analogical reasoning demonstrably occurs, and argue that it crucially depends on lawyers' expertise in the legal domain, which allows them to rapidly and unconsciously assess relevant similarity between two cases (Schauer & Spellman, 2017). In so doing, Schauer

[2] Sunstein may agree: he states that "[w]ithout analogies, relevant principles often cannot be described in advance except at an uninformatively high and crude level of generality" (Sunstein, 1993). That is, an analogy may be necessary to discern whether a rule actually applies to the facts of a case.

and Spellman implicate a factor critical to everyday analogical reasoning. Weinreb also discusses the role of expertise, along with other factors that govern analogy in everyday reasoning, including commonsense reasoning, rich representations, and multiple representations, which are described in more detail below.

B. MODELING LEGAL REASONING

This research builds upon several facets of the AI & Law literature. The most closely related prior work is GREBE (Branting L. K., 1991), a system that used a combination of first-principles and case-based reasoning by analogy to evaluate worker-compensation claims. Given a description of a set of case facts (in propositional logic), GREBE used first principles to derive facts known to be relevant to the analysis of workers' compensation claims. It then used both backchaining and its own A*-based implementation of Structure Mapping (Gentner, 1983) to create mappings to all past cases (20 precedents converted to 35 precedent constituents representing different case components); it output an answer for the case at bar by applying the answer from the most analogous prior case (GREBE defined "most analogous" to mean "with the greatest portion of the current and past cases aligned"). Precedential reasoning in GREBE was thus used to derive intermediate propositions in the larger rule-based reasoning system: the ability to derive a fact by analogy was treated as identical to the ability to derive that fact through backchaining. GREBE was tested on 20 hypothetical cases. The current research moves beyond GREBE in several ways. First, my work starts from natural-language inputs rather than hand-coding cases into propositional logic representations, and does not involve hand-encoded rules. My work also focuses on modeling and generating legal arguments (i.e., a lawyer's reasoning) as well as objective legal analysis (i.e., a judge's reasoning). Most significantly, because my system does not involve writing any rules by hand, it is at least theoretically applicable to any common-law domain without needing an expert to specify legally-relevant rules or information (whereas GREBE was specialized on workers' compensation claims, and would require other hand-coded rules to reason in other domains).

The most productive and active area of AI & Law research relevant to the present research is the HYPO family of legal reasoning, case-based reasoning, and argumentation ((Ashley K. D., 1991); see also CATO (Ashley K. D., 1999), BankXX (Rissland, Skalak, & Friedman, 1996), SMILE+IBP (Ashley & Bruninghaus, 2006)). The HYPO family of modeling techniques reasons in domain-specific areas; many of these systems, for example, are specialized for reasoning about trade-secrets cases (although BankXX is specialized to reason about the applicability of tax-code deductions). HYPO-style reasoners operate over *factors*, which are legally relevant concepts (essentially labeled fact-patterns) that may or may not be present in a given situation (some early versions of these systems used *dimensions*, which are valenced factors: they indicate not only whether a factor is present in a case or not, but the extent to which that factor is present). Early versions of these systems had to have both a human-readable squib describing the case and a hand-encoded factor representation for reasoning by the AI system, but with SMILE+IBP, Ashley & Bruninghaus extended the systems to recognize factors from the texts describing cases. The factors still have to be defined by hand by humans, but the system is now able to recognize whether those factors are present in individual cases or not.

Given a scenario, the HYPO family of algorithms operates in three steps. First, an analogous case is retrieved and proposed as an analogue: all cases in the case library labeled as sharing factors with the case under consideration are retrieved and the cases that share the most factors with the case under consideration are selected as most on point. The shared factors are used to construct an argument that the case under consideration should have the same outcome as the retrieved case. Then the system argues against that retrieved case, by distinguishing it from the case being considered, by proposing a different case as the appropriate analogue, or by reinterpreting the mapping. Finally, the system responds to and resolves this counterargument. By incorporating the counterarguments into the system's reasoning, the system can arrive at more reliable, more thoroughly argued solutions.

It is worth noting here that the representations of factors and dimensions have become the primary representational formalism of the AI & Law literature because of the simplicity of implementing such systems and the fact that using them allows researchers to sidestep the problem of having their system determine which generalized fact patterns represent important legal information.[3] That is, humans define what kinds of intermediate legal concepts are relevant to deciding a case; the system need not itself make that determination (as noted above, Ashley & Bruninghaus developed a technique to automatically determine which factors were present in a case description, but the factors themselves were still provided by humans). My research eschews factor-based representations, which depend on humans defining what sorts of legal concepts the systems might recognize. Instead, my research uses commonsense concepts applicable across a variety of human situations; legally relevant sets of facts (the "factors") emerge from the shared information discerned across cases. The generalization of fact patterns in my research can thus be seen as a means to learn not only *which* factors are present in a case, but the *content* of those factors.

[3] That said, there has also been extensive work on how to represent "open-textured terms," e.g., legal terms that bear multiple meanings and are left intentionally vague so as to be applicable in multiple contexts; see, e.g., (Rissland & Skalak, 1989), (Bench-Capon, 1993).

My work also builds upon extensive work wherein cases are represented as formal logical statements that, in combination with the outcomes of the case, encode rules and arguments. For example, (Horty, 2011) developed a logical formalism for weighting rules extracted from precedent cases. The reasoning in Horty's system was all done in formal logic rather than using relational analogies that leveraged shared structures, but represented an important step in reasoning about cases using rules extracted from and weighted according to their use in a case base. Similarly, (Verheij, 2017) uses case models, or a selection of cases that are collectively logically consistent, different, and mutually compatible, such that the cases together encode the logical rules defining the workings of an overall system. These cases can then be applied either through the rules they encode, or through what Verheij calls analogies, which essentially involve applying a case's outcome on the basis of one or more shared statements in the case's logical representation. In this way the logical rules defining a system can be used without having access to the entire rule set. Horty and Verheij's systems are only two among several that use formal analogy over cases represented in a classical formalized logic;[4] while the general approaches of such systems offer useful guidance for how to manage and leverage case bases, they assume a formalism that is not yet easily extracted from natural-language descriptions of cases. In using explicit representations that can be extracted automatically from language, the present research expands the lessons and results of these past systems into broader areas. Furthermore, using an analogical reasoning system that leverages shared structures (i.e., relationships between events and entities) rather than only shared features results in a richer representation of legal events, and a more flexible (if less inferentially sound) tool for performing precedential reasoning.

[4] There is more research still that eschews formal logic entirely and simply lists cases as a collection of features disconnected from each other, e.g., (Walton, 2010; Kannai, Schild, & Zeleznikow, 2014).

Branting and colleagues developed a deep learning system to both predict and explain legal decisions, and trained and tested it on just over 16,000 World Intellectual Property Organization domain name dispute cases (Branting, et al., 2019). They manually tagged individual sentences in a small number of cases' statements of fact with factor-like representations, then trained word- and sentence-embedding systems to propagate those tags to the other cases in the dataset. The idea was that a system trained on text and tags would not only be able to perform prediction, but could use those tags to identify to a user what in the case had led to the prediction. They found that prediction on a system trained on raw text was more effective than one trained on tagged text. Though they did not evaluate the explanatory capabilities of their system, they argue that being able to explain a legal prediction is critical for any users of such predictive systems.

Finally, the AI & Law literature has long been divided into work aimed at predicting outcomes in cases (see, e.g. (Bruninghaus & Ashley, 2003; Horty, 2011)) and work aimed at generating arguments for one or both sides in a case, without regard for what the "objective" outcome for the case should be (see, e.g., (Aleven & Ashley, 1997; Prakken, Wyner, Bench-Capon, & Atikinson, 2013; Al-Abdulkarim, et al., 2019)). The present research extends and unites both lines of research by developing algorithms usable for both prediction and argumentation.

C. STRUCTURE MAPPING THEORY AND THE COGNITIVE SCIENCE OF ANALOGY

This research is grounded in Cognitive Science research on reasoning by analogy, particularly Structure-Mapping Theory (Gentner, 1983; Gentner & Markman, 1997; Gentner & Smith, 2013).[5] Structure-Mapping Theory posits that reasoning by analogy involves aligning relational cases—cases where the relationships between the entities involved are understood—and

[5] While Structure-Mapping Theory is not the only psychological account of analogical reasoning in humans—see, e.g., multiconstraint theory (Holyoak & Thagard, 1997)—it is the theory within which my research is conducted.

using the alignment to leverage what is understood about one case to reason about the other. The alignment of relations, rather than aligning only attributes, is what drives reasoning by analogy (that is, while attributes may be aligned, the relational alignment is the key to the analogy). Structure-Mapping research has demonstrated that such reasoning is guided by constraints, including that items in one case can map to at most one item in another, and alignments that map deeper shared structures are preferred over those that do not.

Despite the fact that effective reasoning by analogy is driven by deep, shared relational structures, retrieving such cases can be difficult for humans. Retrieval of cases from memory is guided by surface similarity, that is, by shared features and attributes of entities (like entity type) rather than the relationships between them. (Gentner & Landers, 1985; Ross B. H., 1987; Catrambone & Holyoak, 1989; Gentner, Rattermann, & Forbus, 1993; Trench & Minervino, 2015). Because lawyers search through databases to find relevant precedent cases as well as through their own memories (if they are familiar with an area of the law), this facet of human cognition is less applicable to legal case retrieval than to human memory retrieval, but it has informed the design of the case-library retrieval tool my research relies on (see below). It also helps explain why experts are better at retrieving cases useful for analogical reasoning than novices. (Novick, 1988; Rottman, Gentner, & Goldwater, 2012). Because experts have richer relational understandings in their domains of expertise (i.e., they understand the mechanisms underlying the operation of those areas), they have the kinds of rich relational representations that are useful for analogical reasoning. This same richness of representation helps explain why they are better at retrieval: being familiar with a domain, they can label those deep relational structures as emblematic of a certain kind of pattern within the domain, collapsing that relational structure into a feature for the purposes of retrieval (a “structural feature”) (Novick, 1988).

My own research does not focus on retrieval, and will eschew labeling situations with expert *legal* concepts, as in Ashley's work and its progeny. Nonetheless, the lesson of expertise is clear: effective reasoning by analogy requires a rich understanding of the domain about which reasoning occurs, which for my purposes (modeling Tort law) is everyday situations involving people wronging each other. My research therefore depended on ensuring that the events in the tort cases in question are understood at the level of an intelligent adult human (for example, understanding that if someone points a gun at another, a threat has occurred, and likely materially differs from the same situation but with an obviously fake squirt gun). The particular techniques used to ensure this occurred are discussed in Chapter 4:.

Finally, structure-mapping is easier and more effective when reasoning with schemas than reasoning from individual cases (Gick & Holyoak, 1983; Ross & Kennedy, 1990; Bowdle & Gentner, 2005). This is because such schemas emphasize shared structure and deemphasize surface features, which can serve as distractors to analogical reasoning. Furthermore, constructing analogical schemas more effectively occurs through comparison of similar cases than through focused analysis of what "drives" a single case. (Gick & Holyoak, 1983; Gentner, Loewenstein, Thompson, & Forbus, 2009). Constructing these schemas across dissimilar cases can be difficult for people without some signal prompting them to compare the cases (Catrambone & Holyoak, 1989). This again helps explain the supremacy of experts in reasoning by analogy: having seen more cases in a domain, experts can hone in effectively on the shared structures and ignore the distracting features which, they understand, are irrelevant to what happened in any given case. But it is also an explanation of how expertise is acquired, because people with no expertise can be made to rely on relational similarity within a given domain by being directed to actively compare examples illustrating that domain.

Note that though the structure-mapping process is the same, reasoning relative to a schema might not be considered an “analogy”, which involves reasoning between two cases; a comparison between an abstract categorical schema and a grounded case is called a *vertical alignment* (Bowdle & Gentner, 2005). That said, this thesis will generally use the terms “analogical reasoning” and “analogy” to refer to the structure-mapping process and its outputs.

D. THE STRUCTURE-MAPPING ENGINE, COMPANIONS NATURAL LANGUAGE UNDERSTANDING, AND THE KNOWLEDGE BASE

This research is implemented using the Companions Cognitive Architecture (Forbus & Hinrichs, 2017). It includes the Structure Mapping Engine, an ontology and knowledge base derived from Cyc— a generation-long effort to encode human knowledge in a propositional-logic-based knowledge base (Lenat, 1995)—and the Companions Natural Language Understanding system. The default Companions knowledge base (NextKB) that this research was performed with contains over 18,000 concepts and over 1000 types of relations, constrained by nearly a million facts. Many are drawn from OpenCyc (the freely-distributed subset of the Cyc knowledge base), but the knowledge base has been supplemented with semantic and lexical information and support for qualitative and analogical reasoning and learning. Knowledge is partitioned into over one thousand microtheories (contexts which scope the truth value of facts), which can be linked via inheritance relationships to form logical environments to support and control reasoning. Microtheories allow facts that are true in different contexts—but which would cause a contradiction if reasoned about together—to coexist in a knowledge base and be reasoned about in their appropriate contexts. For example, in the RealWorldMT Luke Skywalker is a fictional character and magic is not real (except as subterfuge). But in the StarWarsMT Luke Skywalker is a person who can lift objects with his mind. Both these microtheories could inherit from a third, like FamilialRelationshipsMT. This microtheory inheritance would allow one system reasoning in RealWorldMT about Princes Harry and William and another system reasoning in StarWarsMT about Luke and Leia to determine, e.g., that each sibling pair must have the same set of parents, without polluting the RealWorldMT with Star Wars facts, or vice versa. Finally, using representations partially derived from OpenCyc enables leveraging the several person-centuries of work that has gone into its development and reduces the risk of tailorability, as does using natural language inputs. NextKB’s ontology has the desired characteristics of a legal ontology described in (Ashley K. D., 2009) and further elaborated in (Atkinson & Bench-Capon, 2019).

Analogy is an important reasoning and decision-making tool. The Structure-Mapping Engine (SME, (Forbus, Ferguson, Lovett, & Gentner, 2017)), one of the two core reasoning tools underpinning the present research, is a computational model of analogy and similarity based on Gentner’s structure mapping theory (Gentner, 1983). SME operates over structured, relational cases: cases encoded in propositional logic where entities and events are described and connected to each other using concepts and relationships that are defined within a larger ontology and used consistently across cases. SME takes in two such cases (the *base* or *source*, from which reasoning proceeds, and the *target*, the case about which reasoning occurs) and computes up to three mappings between them. A mapping includes correspondences between the cases, candidate inferences suggested by it, and a similarity score. If a candidate inference involves an entity not in the other case, that entity is hypothesized as a skolem.

SME mappings are computed consistently with constraints derived from the psychological research described in the previous section. First, in the mapping output by the system, any given entity in a case can map to at most one entity in another case (the *1:1 mapping* constraint). Second, if a statement in one case is mapped to a statement in another, its arguments must also map to each other (the *parallel connectivity* constraint). For example, if `(hits Alex1 Bill2)`—"Alex hits Bill"—in the base is mapped to `(hits Carl3 Dave4)` in the target, then `Alex1` must map to `Carl3` and `Bill2` must map to `Dave4`. Third, identical matches between relations are preferred, but non-identical relations can be mapped if doing so supports mapping larger structures (the *tiered identicality* constraint). For example, if the base contains `(insults Edna1 Francis2)` and the target contains `(attacks Greta3 Helen4)`,[6] the statements will not align by default, but can be aligned if doing so lets `(causes (insults Edna1 Francis2) (isUpset Francis2))` map to `(causes (attacks Greta3 Helen4) (isUpset Helen4))`. The final constraint is *systematicity*, which says that mappings that align larger systems of relationships between the base and target are preferred over ones that align more disconnected systems of relationships.

SME proceeds in three phases. In the first phase, the algorithm discovers all possible local identity matches between expressions: shared identical relations and entity types are matched to each other.[7] If relations match to each other, their arguments will be proposed as potential matches as well, even if those arguments would otherwise not be matched at this phase. No consistency constraints are enforced at this phase; these *match hypotheses* instead provide the initial ingredients from which mappings will be generated. In phase two, SME assembles consistent match hypotheses into *kernels*: it assembles shared linked statements into chunks that do not violate the 1:1 mapping and parallel connectivity constraints. After phase two, the kernels are themselves internally consistent, but different kernels may be inconsistent with each other. In phase three, kernels are assembled together: starting with the largest kernels (i.e., those that map the most

[6] I'm modeling Torts, here; people in those cases don't tend to treat each other nicely.

[7] My systems rely on neo-Davidsonian event representations, where events are reified as entities to support compositional reasoning, which allows identical event types to align to each other; this means `(hits Alex Bill)` would be represented as `(isa hit123 HittingEvent)`, `(performedBy hit123 Alex)`, etc., and could map to other HittingEvents in the target.

shared structure), the algorithm takes the next largest consistent kernel (i.e., that does not violate mapping constraints) and adds it to the mapping; it continues until no more consistent kernels remain. It then starts over with the largest unmapped kernel. At the end of phase three, candidate inferences are projected from one case to the other. (For the most recent paper that focuses comprehensively on SME, see Forbus et al., 2017.)

MAC/FAC is a model of analogical retrieval inspired by human cognition that combines cheap, feature-based case retrieval with an analogical comparison to the retrieved cases (Forbus, Gentner, & Law, 1995). MAC/FAC retrieves cases that may be helpful for analogical reasoning from a case library, without relying on any indexing scheme. It takes in a probe case like those used by SME as well as a case library of other such cases. MAC/FAC efficiently generates remindings (SME mappings) for the probe case with the most similar case retrieved from the case library. MAC/FAC operates in two phases, the first of which is driven by feature similarity between the probe case and the cases in memory. This means MAC/FAC is not guaranteed to retrieve the case most analogous to a probe. As a model of everyday human cognition, this is a feature, not a bug, since humans display the same retrieval patterns as MAC/FAC (retrieving cases that share features over those that share only relations). But this pattern is less true of case retrieval in a legal setting.

Given a probe case, MAC/FAC retrieves up to three cases and returns up to three SME mappings from those cases. The first stage, MAC, first computes content vectors (CVs) for a case. CVs compactly encode the attributes (entity types) and relations in a case. CVs are stored for cases in memory to avoid recomputing them, and are normalized according to the size of the CV. MAC then performs dot-products between the CVs of the probe case and of each case in memory. This serves as a fast, coarse estimate of the overlap between the two cases. Up to the top three cases (if the scores are similar enough) are then passed to FAC. FAC generates SME mappings between the probe case and each case in memory, again returning up to the top three. MAC/FAC's operation method is not core to the claims of this thesis—indeed, improving case retrieval given the attributes of the Dataset is a discussed area of future work.

SAGE (Kandaswamy & Forbus, 2012) is a model of analogical generalization built on SME and MAC/FAC. SAGE is the other reasoning tool core to my system. Given a new case, SAGE uses MAC/FAC to retrieve a similar case or generalization. If sufficiently similar, SAGE uses SME to generalize the cases together (future cases can be further assimilated). If the case is too dissimilar, it is stored to potentially generalize with future cases. These generalizations can be used as cases for further SME comparisons. Rather than keep only facts common to all generalized cases, SAGE generalizations are a joint distribution over the facts of all constituent cases. Each fact is stored in the generalization together with its probability, i.e., the proportion of cases in that generalization that contain it. Facts whose probability falls below a preset threshold are removed from the generalization. Thus the generalization can maintain information about which facts are likely, not only which are universal. To illustrate, consider a generalization made of three cases that describe dogs: a Golden Retriever, a yellow Labrador, and a Dalmatian. The generalization will have the fact that a dog has 4 legs with probability 1.0 and the fact that it has yellow fur with a probability of 0.67.

The system generates the CycL-style representations for cases by processing simplified descriptions of cases through the Companions Natural Language Understanding system (CNLU, (Tomai & Forbus, 2009)). CNLU produces hierarchical parse trees using Allen's bottom-up chart parser (Allen, 1994). At the leaf nodes of the trees (individual words or compound phrases), subcategorization frames are retrieved and used to generate choice sets. Interpretations are formed by automatically selecting consistent sets of choices (Barbella & Forbus, 2015). Coreference resolution merges different references to the same underlying token.

CNLU operates most effectively over a simplified English syntax, roughly that of elementary school materials. Researchers using CNLU often use simplified syntax to focus on semantic breadth, the range of ideas that can be expressed in the underlying representation, over syntactic breadth, the range of surface forms that can be processed. CNLU uses Discourse Representation Theory (Kamp & Reyle, 1993), implemented via microtheory inheritance, to construct a full semantic description of sentence content. This allows the system to handle negation, implication, quantification, and counterfactuals, using nested discourse representation structures (DRSes).[8] Once language processing is complete, these DRSes are converted to standard CycL representations and scoped by microtheories.

CNLU uses lexical and semantic knowledge from NextKB. The lexical information – which includes parser information – was semi-automatically extracted from a public domain edition of Webster's dictionary and augmented with current vocabulary. Semantic translations derived from FrameNet (Ruppenhofer J. , Ellsworth, Schwarzer-Petruck, Johnson, & Scheffczyk, 2016) map the lexicon to the NextKB ontology. Each word, for each part of speech, has one or more semantic translations to express possible meanings. Complements in the syntactic analysis

Figure 2.1. "Dave eats Ice Cream" in traditional predicate logic (L) and in neo-Davidsonian event representation (R).

(isa treat123 IceCream) (eats Dave treat123)	(isa treat123 IceCream) (isa eat456 EatingEvent) (doneBy eat456 Dave) (objectConsumed eat456 treat123)

[8] Discourse Representation Structures scope facts like microtheories do, but according to the facts' role in a discourse. For example, the sentence "Jane said that Keith did not breakfast" would yield four DRSes: The first contains the facts that Keith *did* eat breakfast, the second wraps the first in a negation (indicating that Keith did *not* eat breakfast), the third states that Jane says something, and the fourth links Jane's speech to its content.

are mapped to role relations. Thus the system recognizes that “eats” is an instance of Eat-TheWord which can refer to an instance of EatingEvent, and that “Dave eats ice cream” makes Dave the event’s performer. CNLU generally uses neo-Davidsonian event representation: events are reified and related to other information (e.g. actors, location, etc.) to allow for greater composability (Figure 2.1). Neo-Davidsonian representations allow additional information to be added incrementally by relating it to the underlying event, rather than needing to create a new statement. Part of CNLU’s interpretation process can involve *Narrative Functions* (McFate, Forbus, & Hinrichs, 2014), which are abductive explanations for a statement. They are detected by rules that operate across choice sets to construct explanations for a parse, making abductive assumptions as needed.

This work is also a continuation of MoralDM (Deghahni, Tomai, Forbus, & Klenk, 2008; Dehghani, Sachdeva, Ekhtiari, Gentner, & Forbus, 2009; Blass & Forbus, 2015). MoralDM solved moral dilemmas taken from psychology research studies. MoralDM used two reasoning systems: first-principles reasoning, based on established psychological moral reasoning principles including *protected values* (inflexible moral values, the application of which constitutes obeisance of moral rules) and the *principle of double effect* (a moral principle), and analogical reasoning against solved cases. The solutions to moral dilemmas depend on the particular moral norms of the person making the decision and are therefore inherently subjective, yet within particular cultures and groups overall trends do emerge; MoralDM took the majority decision of participants in the original studies it was modeling as being correct.

MoralDM uses a combination of first-principles rules and analogy to detect protected values. MoralDM starts with some first-principles reasoning to establish protected values and other facts to prime comparisons, then uses analogy over cases for which it knows the right answer. It uses MAC/FAC to retrieve cases, SAGE to construct generalizations from cases, and a domain-independent consistency check to evaluate whether a retrieved case is a reasonable one with which to reason about the new case. That is, it tests to see if a candidate inference is already known to be false in the target, and if so, ignores the match. MoralDM was able to successfully solve moral dilemmas by leveraging the power of generalization and analogical inference, built atop a base layer of first-principles reasoning.

This describes the research background underpinning the thesis. But before turning to the AI research I conducted, it is important to discuss why and how I believe AI research modeling legal reasoning and argumentation should be done. The legal system is a core pillar of a functioning society—"move fast and break things" will not do.

CHAPTER 3: HOW TO APPROACH MODELING LEGAL REASONING

This chapter explains in more depth the goals of the system, and how those goals informed why certain engineering approaches were taken and others rejected. It proposes a new—to the legal academy, at least—role for analogy in legal reasoning, namely, as the mechanism by prior cases are compared *to each other* in order to discern the rules governing those cases. It discusses potential roles for legal automated reasoning systems, what properties those systems should have (particularly transparency and explainability), the utility of cognitive modeling, and the desired scalability of the PARALegal system.

A. GOALS

My goal in developing PARALegal is not to replace legal decision-makers like judges, legal advocates like lawyers, or support staff like judicial clerks or paralegals (despite the pithy name). In fact, I believe that *no* automated legal reasoner should replace human decision-making, and these reasons should be discussed before moving to the goals of the thesis.

i. Why Not to Automate Legal Reasoning, Argument, and Decision-Making

There is already ample evidence that where legal automation has taken place, it has produced negative results on the whole (Citron, 2008), although the success of systems like DoNotPay suggest perhaps a circumspect role for legal automation in contesting bureaucratic enforcement actions. But automated legal reasoners should not, even in theory, replace human legal actors, for three theoretical and one pragmatic reasons: the impossibility of fidelity, the inherent humanity involved in determining what constitutes “justice”, the at-times-necessary arbitrariness of legal judgment of facts, and how AI algorithms get deployed in the real world. Each of these is fundamentally about procedural justice, the justice that inheres in procedures, as opposed to rules and outcomes. Each is discussed briefly before turning to the goals of developing the system.

Faithful Replication of Legal Decision-Making is Impossible

Either an AI judicial decision-maker would be an explicit alteration of the existing system, or it would attempt to be a *model* of it. If it is a model, then it is by definition an imperfect one, because models require simplification (O'Neil, 2016). In my Article *Observing the Effects of Automating the Judicial System with Behavioral Equivalence* (Blass J. A., 2022) I unpack this point from a technical perspective to argue *why* it is theoretically impossible to replace a component of the legal system and leave the larger system entirely unchanged in its operations. The curious reader is directed to the Article, which at 28,000 words goes into more depth than would be appropriate for this thesis, but the argument can be summarized as follows.

Computer scientists (and other engineers) define whether two systems “work the same way” using a concept called “behavioral equivalence”. Behavioral equivalence looks at the specific behaviors of two systems: two systems “work the same way” if they display equivalent behaviors. But what counts as a behavior? It turns out that behaviors are always and necessarily defined relative to some observer: behaviors are equivalent if they are *observed* to be equivalent (Felleisen, 2009). The observer should be constructed to only observe that part of the behavior that actually matters to the performance of the system. If some system performs addition using base-10 vs. base-2 (binary) numerals as its internal system, that presumably will not matter to any observer who is interested only in the output of the system. But what if one system performs addition by adding in columns from right to left, and another does so by incrementing one of the addends by one, as many times as the other addend? The former system will perform addition in roughly as many operations as there are orders of magnitude of the numbers being added; the latter will require as many operations as one of the numbers itself. Then, the choice of observer determines what counts as a side effect: an observer that only looks at the outcomes will not differentiate between a system that is computationally efficient and one that would take a significant amount of time simply to add two numbers together. Instead, an observer that looks at the internals of the system should be chosen for this example if the efficiency of various systems is to be observed.

The choice of observer thus determines whether the differences between two systems are perceived, and the two systems are therefore distinguishable. To a child who wakes up in the morning and finds clean tableware in the cupboard, a mechanical dishwasher is equivalent to a dishwashing parent; not so to the parent. And this leads to the second critical point about observational equivalence in the legal context: to the omniscient observer who sees every difference, no two systems can be equivalent (Morris, 1968), even multiple copies of the same system. Even the same computation performed by the same system at two different times can be distinguished by an observer that observes the time at which a process is run; even two copies of the same system that perform the same computation at the same time can be distinguished by an observer that observes the hardware the system is running on. This is not to say that one system is good or bad, only that if everything that could possibly be observed is observed, no two different systems will be *equivalent*. And that leads to the final point about observational equivalence: that the collection of all possible observers is equivalent to the omniscient observer, because the salient question is not who the observer is, but what behaviors are being observed. If everything about a system is observed, any change to that system will be detected, and the new system will not be understood as being the same as the old one.

What does any of this have to do with the law and automating legal decision-makers? It has to do with *procedural justice*, the justice that inheres in a system's procedures, not just its outcomes (Rawls, 1999; Nozick, 1974). People can be more accepting of an outcome they do not like if they feel they got a fair hearing than they are of their preferred outcomes distributed arbitrarily (Lind & Tyler, 1988). Not just legal outcomes but legal processes matter, and not only to participants in the legal system, but to society at large. Thus the legal system and its processes appear to present a case where everything about some system is observed, and therefore no other process can be equivalent to it. Litigants and lawyers observe what happens in the courtroom and to their case. Judges and clerks observe what happens in chambers. But the legal system and what happens within a courtroom matters to more than just those who come into the courthouse. Even if a faithful automated replica of a trial process could be engineered, the society within which that trial process is embedded would not perceive the replica as being the same as the original system. What happens inside courtrooms gets broadcast to the rest of society not only by the litigants, but by the media. Not only case outcomes but the rules of decision that resolved them are important to the public. Those rules of decision are scrutinized and modified not only by higher courts, but by government entities, which are responsive and accountable to the people. Interest groups are enormously varied and care about the extent to which their interests and values are advanced (or not) by legal rules and processes. And potentially any part of a legal process can affect one of those interests.

Think of it this way: individual cases (and the lawyers and litigants involved) are like blood coursing through a body. To the blood, a dialysis machine might work as well as a kidney; not so to the body whose blood is being cleansed. So even if some automated process faithfully replicated some aspect of legal reasoning or decision-making, it could not be a faithful replica to *everyone*, because some observer would notice that which had changed. And that argues for caution when automating legal processes because there is no clear principle by which to carve off parts of the legal system and say "it would not matter if this changed, so our observer of behavioral equivalence can ignore it." On a case by case basis it might be possible to determine that some change in observed behavior can be safely ignored (e.g., someone might notice if the coffee brand used in the jury pool room was changed), but the law of unintended consequences suggests being cautious when making changes and being confident that one understands what exact changes will result.

But of course, there are recognizable problems with the legal system that should be fixed, in which case the observed difference in the system would be the point of making the change. And certainly one can imagine situations in which using an automated system might lead to improvement over the status quo. Nonetheless, there are good reasons to be cautious when attempting to automate some aspect of the legal system. Even if one can overcome the potentially large hurdle of translating policy into code—an inherent process of translation, and therefore distortion (Citron, 2008)—some aspects of judicial decision-making simply should not be turned over to computers. For example, one of the judiciary's roles is to interpret ambiguous law, which in practice can mean deciding what the law should be. As the next section argues, AI systems are ill-suited to this task.

JUSTICE AS A (SOMETIMES) INHERENTLY HUMAN CONCEPT

Procedural justice, as previously noted, represents the idea that justice inheres in procedures themselves, and not only in the decision rules applied to cases, and what outcomes are derived from those rules. But the rules and outcomes, of course, matter as well. And in the legal system, the rules are determined by humans: either by groups of humans organized as a legislature and empowered to come together to create the rules, or by individual judges—or panels of judges on a reviewing court—crafting common law rules. And when a judge is trying to craft a rule for a particular case (that is not entirely specified by inescapable binding precedent), she is trying to craft the *right* rule, based not only on precedents that underspecify what the rules in hard cases should be, but also based on what is right, just, and fair for the community that will have to live under that rule (Dworkin, 1986). This ability to "know what to do when you don't know what to do" is, as yet, a fundamentally human ability. I could stop there, because it is enough simply to say that a skill fundamental to legal judgment is, for now, in the realm of science-fiction, but the ability to know what to do when you don't know what to do may not be science-fiction for long. AI systems might be developed that have overriding norms in them, and those norms might conflict in interesting ways that lead to interesting emergent behaviors (Olson, 2021). Or they might be trained on a sufficiently enormous and varied background dataset that they display interesting and unforeseen behavior when confronted with new prompts (Ramesh, Dhariwal, Nichol, Chu, &

Chen, 2022). But still firmly in the realm of science fiction are AI systems that we can guarantee will come to the judicially *right* decision.

There is no guarantee, of course, that a human judge will come to the right decision either. But there are good reasons to assume that, in general, human judges will come to *a* right decision, because that is what procedural justice is all about: society has delegated to the human judge the power to make the right decision, and we give the human judge the benefit of the doubt that they will do their very best to achieve justice in a given case. We accept that laws will be at least somewhat indeterminate, and that judges may engage in "permissible disagreement" about what justice demands from an interpretation of a given law in a given case (Gowder, 2020; Fallon, 2005). We provide legal mechanisms to deal with cases in which a rule has been wrongly crafted or justice has not been served in a case, including review by higher courts, legislatures changing decision rules, and impeachment of judges (or voting in new judges in States where judges are elected). We might disagree with a decision, but unless we feel as though the decision itself is somehow illegitimate (for not being grounded in law, for example), we will still accept it (Strauss, 2005).

We could of course collectively decide to empower some AI system to play the role of a judge, and that system's decisions would be legitimate, provided they acted within the law and were accepted by the society that empowered them (Fallon, 2005). I accept that a sufficient majority of society at large might one day disagree with my personal conviction that it is precisely the fact that a rule was crafted or that a case was decided by a human doing their best that confers legitimacy on the system and allows people to accept the outcomes of the case.[9] That future society

[9] At the very least, not having a human decision-maker to at least review automated judgments would not be legitimate under our current system, where our national and State Constitutions specify how (human) judges are

may well empower an AI legal decisionmaker whose decisions I and others who share my concerns will have to accept as legitimate. Perhaps some AI system might one day be developed that can justly dispose of relatively straightforward cases. Nonetheless, there is one judicial function that is, at least sometimes, inherently human: determining what the rules governing cases should be.

Examine the relatively rare judicial instance where it is genuinely unclear what rules govern some incident, or what those rules mean, and the judge must craft one in order to dispose of the case. The judge knows of the rules governing her own conduct, in general and with regards to crafting the new rule, and she knows about other rules governing various aspects of the world. She also has a set of *premises*: what she knows or assumes to be true, both about the case and about the world outside the case (including the existence of other legal rules). Some of these premises were themselves at one point inferences drawn from prior premises and rules: for example, that matters on which the pleadings agree can be treated as facts. The judge's task is to use these premises and rules to make inferences, in this instance, what should be the rule that will dispose of the case.

In the situation where the rule is unclear and needs to be crafted, or might be reconsidered, or does not dispose entirely of the case and leaves the adjudicator to decide what is "fair" within the blank space,[10] some of the premises involved in settling on a rule will sometimes be overriding norms about what is "right." When this happens, the judge is operating using premises that involve

---

selected and confirmed. There is also evidence that humans intuitively see automated *legal* judgment as being less fair than human-made legal judgment (Chen, Stremitzer, & Tobia, 2022). Those findings run counter to the extensive study of "automation bias," wherein an automated decision-maker is seen as more objective and trustworthy than a human one (Alberdi, Ayton, Povyakalo, & Strigini, 2005; Lee & See, 2004; Bussone, Stumpf, & O'Sullivan, 2015; Lustig, et al., 2016).

[10] This relates to another practical objection: that systems trained on or designed for existing law cannot understand when the law is unjust and must be changed. Given that the current discussion is on the theoretical objections, however, let us assume that such an AI system exists.

value judgments about which reasonable people may disagree.[11] For example, some people prize individual autonomy over collective harmony, while others believe individual autonomy should face strict limits in the face of collective need. Some people think government oversight is intrusive, while others believe it is protective. People can examine why they feel certain ways, but eventually those feelings ground out in just that: a *feeling* of what is right, or what really matters. Fairness is an ethical consideration, not a technical one (O'Neil, 2016). This is true regardless of whether the judge is relying on her personal convictions of right and wrong, or doing her level best to discern what "political morality" best justifies our existing body of law and therefore points to what moral considerations should inform the new decision (Dworkin, 1978). Eventually the task of crafting a rule will require using a premise that is a value judgment not universally agreed upon.

Before I address why this might even be a problem, the objection could be raised that perhaps this will not be an issue for AI systems, because a sufficiently sophisticated AI system might be able to start with universally accepted premises and use those to reason its way to any intermediate value judgments that it needs to determine what rule ought to govern a case. But philosophers have tried throughout human history to derive ethical systems using pure reason, and while certainly such ethical systems have been developed, they are not universally accepted. Humans have not developed ethical premises that can dispose of every case and that are also widely accepted by society. Perhaps someday a fully generally-intelligent AI will be able to, but the capacities of such a system—one that can solve philosophical problems that have stymied humans—are difficult to imagine, and engaging with it would turn this thesis into a work of

[11] Dworkin argues that every legal and moral question has a right answer, but also acknowledges that mere humans will be unable to derive them in any given generic case (Dworkin, 1986). I do not find Dworkin's argument on this point convincing, but perhaps those who do will read into the future of artificial intelligence the possibility of creating an artificial Hercules, Dworkin's imaginary judge with limitless reason, intelligence, and time.

speculative fiction. Better to assume that an AI system that must rely on a value judgment to craft a rule must do so using an inferential chain that grounds out in values placed in it by humans, values not universally shared by all people. If so, then the concepts of Justice derived from those values is traceable to the humans who placed it in the machine rather than the system itself. In that case, the fact that the AI system has some "value" is inherently arbitrary, in the sense that the system would have different values if it had been programmed by different humans who had implanted in it a different but equally legitimate value. If the value is arbitrary, then the rule derived from it will be as well.

On the other hand, regardless of where the human judge was first exposed to that premise, that value judgment has survived and been shaped by the lifetime of the judge's lived experience. That means that when a human judge grounds out her reasoning in her perspective on what is right—when she makes a policy-based decision—even those who disagree with the conclusion or the premise can understand that she actually believes that she did what is right.[12] When an AI system does so, what should make people think that the AI actually *believes* its premises? The AI system's premises—the norms underlying the ethics that allow it to decide what is fair, if not the ethics themselves—will have been those placed into it by its human creators, and will not be universally accepted.[13] That is, even if some of the AI system's ethical norms are ones it learned

[12] A clear exception is when people perceive judges to have started with their preferred conclusion and backfilled in the reasoning to justify it, but that case does not fall within this discussion of carefully-reasoned opinions. Some readers may see this exception as carving out much of judicial decision-making. When a judge starts with the conclusion and backfills it, the judge relies *only* on values to indicate how the case should come out, not on other sources of legal information and other modes of legal reasoning. If the choice is truly only between partisan judges who are unconstrained except by their personal beliefs and an AI system that has no personal beliefs but can be made to reason consistently and in line with predefined principles, perhaps the AI system would be preferable. That said, I believe this is a false dichotomy, and I assume that human judges in general try their best to come to the right decision based on existing law.

[13] Sunstein argues that judges produce *incompletely theorized agreements*, where they promote secondary values that enjoy greater public agreement than the primary values from which they are derived (Sunstein, 1996). For example, there are a variety of reasons that people might want to protect labor unhindered—e.g., that labor unions

or derived or reasoned its way to, those are necessarily based on prior premises, and eventually those premises must ground out in something that was placed into it by a human. This is not necessarily anti-democratic: perhaps if there does exist some set of non-universal values by which an AI could derive all other values that it would need to settle every case and craft every rule, a legislature could democratically vote to instill such values in an AI judge.[14] The objection is that the resulting AI will produce rules that just as easily could have come out the other way had it been given a different set of value premises, which again makes the application of those values arbitrary. If a value can be changed on a dime, it is not much of a value.

Of course, human judges that rule in cases also eventually ground their arguments out in premises that are not universally shared – why is that not equally arbitrary? One reason is that those premises are the result of the judge's lifetime of lived experience: the simple fact that the premise has survived that experience makes it not arbitrary. A different way of putting this is that the judge's value is not arbitrary so long as it is not arbitrary *to the judge*. The human does not only think her beliefs, she *feels* them. She has the conviction of her beliefs, and I believe that this conviction is a crucial element in perceiving judicially crafted rules as legitimate. Even an AI simulacrum of a judge, trained to replicate an individual judge's values and writing style, would only have a model of that judge's value system. A simplification by definition, such a model will lack the richness and subtlety within that value system that come from experience and that are critical to the determination of how values trade-off with each other. That AI might support the

---

protect worker safety versus that they provide an effective means for political organization— and they can agree to do so more harmoniously if they do not discuss their underlying motivations than if they do. It may be true that some premises might be more widely accepted than others, and that careful value selection can lead some hypothetical AI usage to better represent the overall ethos of the people it judges. But an AI system that supports unions without knowing why is even more suspect than one that can justify its "belief". And regardless, it will always have some values with which some large portion of the people reasonably disagree: plenty of people disapprove of unions.

[14] This could raise interesting issues around judicial supremacy and legislative control over judicial decisions.

judge's decision-making processes, and might even be able to convincingly replicate the judge's decision-making in a variety of cases, but should not be assumed to have the capacity to craft new rules based on those values in the same way that the judge does.

Thus, in at least some situations where judges must craft rules or decide what is fair in the margins of existing rules, their ability to do so *justly* depends on their humanity, because it depends on their having personally-felt convictions. If not grounded in conviction, a judgment call under ambiguity is arbitrary, and arbitrariness is the enemy of justice.[15] Now, in the general case an AI's judgment as to what a rule or outcome should be need not be arbitrary (at least not in every sense of the word), if it is explained by reasons which are themselves grounded in arguments that are grounded in other arguments and reasons etc. AI systems that track this reasoning can provide a clear picture as to why some such system came to the conclusion that it did. (Note that this is not the case in many machine-learning approaches that result in uninspectable decisions and whose "reasoning" is neither easily traced nor reconstructed.) But even the most transparent and explainable reasoner—human or AI—will eventually ground out its reasoning in a series of premises. These premises are often uncontroversial (i.e., the facts are found where the pleadings agree, and through a trial to resolve where they do not; decision rules are provided by precedent and statute), so in the general case we can imagine an AI system that does not behave arbitrarily. This may be true even in many instances where a rule needs to be crafted. But it will not be true in the rare case where judgment requires judges to do that which in the general case they are expected *not* to do: bring their values, borne of their experience, to bear in resolving a case.

[15] Indeed, decisions by both judges and juries have been overturned, and judges censured, in situations where it turns out that the decision-maker used a coin-toss to decide a close issue (The Associated Press, 2007; the New York Times, 1982; The Associated Press, 2000).

If I am correct, then even if we could trust some AI system to correctly dispense justice in the overwhelmingly majority of cases, there will remain some cases for which Justice requires that the decision-maker be a human.[16] That will remain true until a generally-intelligent AI is developed that can reasonably be recognized as a person (not a human, but a conscious being) that has the conviction of their beliefs borne from their experience living in the world for which the judges must craft rules.[17] Such a creature may well exist, perhaps even within our lifetimes, but it is not here yet, and systems that fall short of that capacity cannot legitimately decide on their own what qualifies as just in any given case. Rules and values should not be arbitrary. On the other hand:

OPEN-TEXTURED TERMS AND ARBITRARY FACTS

Unlike rules, the determination of facts in cases is sometimes inherently arbitrary, in part because of the ubiquity of the word "reasonable." "The reasonable person" shows up everywhere as the standard by which the legality of the behavior is to be judged: judging the defamatory nature of some statement, foreseeing losses, interpreting the language of a contract and the circumstances of its formation, interpreting administrative action, reacting to provocations in criminal defense, and more (Gardner, 2015). "Reasonableness" is core to the law of Negligence, which turns on the duty to take reasonable care to avoid harming others. And what counts as reasonable? Ask a jury.

---

[16] Maybe this will turn out to be a minority view, and people will not care where an AI's values come from when it makes decisions based on a value with which they disagree. The work cited in footnote 9 certainly shows that is a risk. Another risk comes from the fact that people have been found to find acceptable actions taken by an automated system that they find immoral and unacceptable when taken by a human. (Guerini, Pianesi, & Stock, 2015). Procedural justice is concerned in part with whether the people subject to a system of justice view its procedures as just, so much of my counterargument would be addressed simply through public acceptance of such a system. But procedural justice is not only concerned with the perception of procedures but with whether they are themselves just, and as I have argued, an AI system that is not generally-intelligent cannot guarantee a non-arbitrary procedure.

[17] I was recently informed that many years ago, when I had just started graduate school, I vehemently argued that automated systems could make better judges than humans. I believe the person who told me this, but it is hard for me to imagine having felt that way, and I do not think I will ever believe it again.

Of course the definition of "reasonable" is not entirely arbitrary, but is more of a "I'll know it when I see it" term. Reasonableness turns not only on subtle variations in facts, but on the legal decision-maker's interpretation of what those facts mean. And reasonableness is far from the only legal term that is open to interpretation: the law is full of so-called open-textured terms that do not clearly specify when they apply. These terms can be used intentionally, to give flexibility and discretion to future decision-makers to make sensible judgments on individual case facts, or they can simply arise from general linguistic ambiguity (Hart H. , 1961). Nonetheless, they exist.

Researchers in AI & Law have studied how to handle and represent open-textured terms (Rissland & Friedman, 1995; Branting L. K., 2003; Bench-Capon & Gordon, 2009). Certainly open-textured terms can be reasoned about in terms of prior examples, and nothing prevents an algorithm from being developed (or trained on prior data) to make case-by-case determinations of whether some legal open-textured term has been satisfied. (My intuition is that the public would find automating the resolution of open-textured terms to be less objectionable than automating legal rulemaking.) Nonetheless, and again for reasons of procedural justice, it is important for open-textured terms like "reasonable" to be resolved by humans rather than machines. Even if a computational system is able to exhaust arguments in favor and against a given interpretation and finds it genuinely ambiguous, providing much more detail than a human could, the final decision should be made by a human: whatever decision the human made will be understood to be grounded in the human's experience, whereas the computer's resolution will rightly be seen as arbitrary. Far from helping the situation, a computer rigorously proving that some term is genuinely ambiguous will exacerbate the arbitrariness of its decision by proving that any resolution of the term is arbitrary regardless of who resolves it. If any legal proceeding ever grounds out in the statement "because I said so," it is important for the thing that said so to be a human that was trying their best to do the right thing.

AI IN THE REAL WORLD

Perhaps the reader is unconvinced. Imagining and working to build an artificial general intelligence that would trounce every criticism I just made is fun and hopeful work, a valid research goal and a project that I have at times felt myself a participant in. In the meantime, there remain serious practical objections to automating judicial decision-making. The practical concern has little to do with whether AI systems *could* improve legal reasoning, and everything to do with whether current AI systems *would* do so. Regardless of whether some AI system might be able to improve on the current system, the systems that will actually be implemented to replace or augment aspects of the current regime, at least in the foreseeable future, are unlikely to do so because of error and overenthusiasm on the part of the humans doing the implementation. This is well-trod ground, and includes issues of humans overly trusting in automated decisions even when they should not (Lustig, et al., 2016), the difficulty (if not impossibility) of creating truly unbiased systems using biased datasets (as historical legal data are) (Gonen & Goldberg, 2019), the tendency of the most popular systems of our day to bloviate and prevaricate (Marcus & Davis, 2020), rule-based systems' tendency to be brittle in the face of new forms of input, and more. But while any given AI system's flaws might be addressed, examining how governmental AI systems have been deployed in the real world is a sobering exercise. Such systems have historically been developed, managed, and usually owned by private entities, but the judicial system is a core *public* social good, and should not be privatized. And to date, automated decision-making systems deployed by governments have a poor performance track record.

Governments hire all sorts of people to directly manage their operations and to interface with the public they serve: lawyers, certainly, but also accountants, economists, engineers, investigators, enforcers, and all manner of human infrastructure. They hire computer programmers, too. But when the government needs a computer system built, they typically contract it out to private vendors (Citron, 2008).[18] These private vendors are then allowed to keep the mechanisms of their systems secret as protected intellectual property and trade secrets (Angwin, Larson, Mattu, & Kirchner, 2016; Pasquale, 2015; O'Neil, 2016). Of course, concerns over the privatization of justice could easily be overcome by having any automated judicial system be developed, managed, and owned by the government. But the historical evidence suggests that that will not happen.

It is also worth noting that, all over the country, government functions have *already* been automated, and they have a terrible track record (Citron, 2008; Calo & Citron, 2021; Kroll, et al., 2017). Calo & Citron review some of the litigation challenging "government agency automation's pathologies in varied arenas, including public benefits, jobs, child-welfare, airline travel, and criminal sentencing," and find that "to the extent that [automated systems] are predictable, it is in their misdirection of government services."

Some elements of government functions may well be automatable: if the file of someone applying for a driver's license contains a flag that prevents them from receiving one, for example, a computer program can observe the flag as well as a human. But the justice system is not and should not be a checklist. And even computable government functions present risks, implemented

[18] A related problem is when corporations gather data on users and shares it with the government (Pasquale, 2015). There are types of information that the government is prohibited from gathering on broad swaths of people, but where corporations are not so constrained. As Pasquale says, "once someone else has gathered that information, little stops the government from buying it, demanding it, or even hacking into it. . . . This is not a 'bug' in our surveillance system, but a 'feature'."

as they are from flawed historical data and by flawed humans.[19] An AI system that replicates past inequities could make things worse by obscuring the basis for a decision (or even that a decision has been made) and by giving those decisions the sheen of computational objectivity. Because algorithms cannot be appealed to directly, and they have been nestled within systems that make them difficult to appeal *from* (O'Neil, 2016), even a system that does no more than replicate existing inequities can serve to further entrench them, moreso than if a human were the decision-maker.

All in all, automating legal decision making appears to be a bad idea. The next section will consider why I nonetheless believe it is a good idea to undertake the research that I have.

ii. Goals of the Present Research

As the prior section argued, using an automated legal reasoner to automate legal decision making – that is, to resolve legal cases – is dangerous and not something I believe humanity should be striving for. However, the risk of an automated legal reasoning system being misused to decide cases does not mean that there are no potentially positive uses of such a system. GPS systems can be used to stalk people, but that does not mean that they do not have significant benefits to the modern world. In the same way, the risk of an automated legal reasoner is not inherent to its existence, but connected to the purposes and tasks for which it is to be used.

The goal of the present research is to develop a legal reasoning system that can eventually be used to help potential litigants (and lawyers) understand (a) what the law is, (b) how the law might apply to a set of facts, and (c) how those facts must change or be recast in order to achieve

[19] As one reader pointed out, "we trust flawed humans all the time to build cars, airplanes, bridges, etc., where failures are typically fatal." True! And of course, the human decision-makers that might be replaced are themselves flawed. But a flawed human decision-maker that makes an error does so on a case-by-case basis, whereas a flaw in an automated system's code may result in errors in *every* case, multiplying the harm caused by the system.

the outcome the party seeks. There are significant unmet legal needs in this country: while criminal defendants are guaranteed lawyers by the Constitution, civil litigants are not, and an enormous number of them proceed *pro se*, that is, representing themselves. A study of state courts' cases in 2012–13 found that while plaintiffs in civil cases are very likely to be represented by counsel (96%), less than half of defendants in civil cases in courts of general jurisdiction have a lawyer (46%) (National Center for State Courts, 2015). While 67% of tort defendants have lawyers, the numbers are much lower for cases involving Property or Contract Law. In only 45% of cases were both parties represented, down from 96% of cases twenty years prior. Another survey of state-level studies estimates that 80–90% of litigants are unrepresented by counsel (Steinberg, 2015).

Those numbers only concern the people who make it into court in the first place. While statistics on the number of people who do not even attempt to enter the legal system due to cost burdens are hard to come by, the Legal Services Corporation, which funds legal services for low-income individuals, reports on the requests for legal services received by the organizations it funds. LSC reports that of the 1.9 million annual requests those organizations receive, half of the requests are turned away due to lack of resources, and half of the cases that are taken end up unresolved (Legal Services Corporation, 2022). They also estimate that low-income Americans only seek legal help for substantial legal issues a quarter of the time, and report that half of those who do not seek legal help do so out of concern for the costs of legal help. Extrapolating from their numbers, there are over four million legal issues[20] facing low-income Americans each year for which they do not seek legal help specifically because of the anticipated cost of doing so.

[20] I say over four million because this estimate extrapolates based on the 1.9 million requests for free legal services each year, which certainly undercounts of the total amount of legal help sought by low-income Americans annually, given that some of those requests go to lawyers other than those underwritten by LSC.

As these numbers demonstrate, there are enormous unmet legal needs, and many people who would benefit from having a lawyer do not have one. While the system presented in this thesis is far from being a lawyer, my goal in performing this research is to bring us a bit closer to the future where anyone who needs legal advice will be able to receive it. This goal is fundamentally to engineer a cheaper lawyer, rather than to create a whole new avenue of legal cognition or to transform how law is practiced or who has access to the legal system. As at least one scholar has pointed out (Gowder, 2018), simply providing cheaper lawyers is unlikely to yield egalitarian results, because cheaper lawyers do little to level a legal playing field (because the lawyers are cheaper for everyone). Nonetheless, I believe that cheaper lawyers bring at least one significant egalitarian result, which is allowing more people onto the playing field at the start.

A system like the one I am proposing should be free for most users, or at most cost what it takes to operate it. This is not to lower the barrier to entry, but to ensure that it is not used to take advantage of people. To see why, consider an analogy to how for-profit colleges predatorily target people with the promise that an education and degree from their institution will significantly improve their prospects for career and future earnings (O'Neil, 2016). The owners of an even moderately effective AI lawyer system would have an incentive to encourage people to sue anyone they can think of. Unlike human lawyers, who must actually work a case, each marginal use of the software would cost its owners nothing, so they would have an incentive to get as many clients as they possibly could.

On the other hand, lawyers and law firms could use a system like the one I am proposing to decide whether to take on a case or not. That could have the opposite result that I hope for, by keeping people out of court and denying them legal representation they might have otherwise had. Also, keeping the system free would also expand the pool of people willing to sue their neighbors frivolously, since it would cost them only court fees to do so. But such suits are likely to be unsuccessful if their neighbor could afford a human lawyer (who, I assume, will long continue to perform better than AI ones). If their neighbors cannot afford a human lawyer, they would have access to the same free AI lawyer they were being sued by, giving an even playing field. Furthermore, genuinely frivolous lawsuits are discouraged by the legal system by sometimes granting legal costs to the winning side, discouraging people from bringing such suits. Regardless, one way to address both the use of the system by lawyers and its use in frivolous lawsuits would be to charge professionals and frequent or repeat users. Making the system costly for those who would use it as a professional decision-support tool or those who would use it to abuse the legal system, while keeping it free for others, will discourage its use for such purposes as well as defray costs for its intended users. As long as there is no financial incentive for the system's owner to slip poisoned promises in peoples' ears, I hope that such a system will result in more good than harm.

The system developed for this thesis is, of course, only an experimental system. But even some polished and perfected future version of this system should never be used to replace a legal decision-maker: the system's legal reasoning process does not support the intake and consideration of legal argumentation, it has no means to craft new rules where appropriate or understand when to depart from prior rules, and there is nothing in its function that addresses the procedural justice concerns that were just described. Because the system cannot distinguish prior precedents, recast case facts, or justify attacking or positing particular facts in a case, it should not be used to replace a legal advocate. It should strictly be used as a support system to help understand what the law is and how it might apply in a case.

B. ENGINEERING PREMISES AND MOTIVATIONAL COMMITMENTS

As I noted, the purpose of this research is to study how to develop a system that accomplishes a particular task. But when developing an AI system to accomplish some task, particularly a task that humans are able to accomplish, it is always appropriate to inquire whether the engineering should be inspired by, or even seek to replicate, the human cognition that accomplishes that task. That's true even when the purpose of the research is not to validate a particular model of human intelligence: because humans are smart and have wide-ranging abilities, it can be wise to take human cognition as a roadmap when engineering AI systems. This is especially true in domains where there is a great deal of evidence of *how* humans accomplish particular tasks. The law may appear to be one such domain, given that legal reasoning is even more inextricably tied to human reasoning capabilities than other demonstrations of intelligence: while other animals can access memories, plan actions, build mental maps, etc., only humans perform legal reasoning.

Nonetheless, this thesis, which seeks to develop an AI model of legal reasoning, decision-making, and argumentation, explicitly eschews any attempt to be a faithful model of human legal reasoning, and does not claim to be such a model. It does not seek to purposely reason in an *in*human way—it does not reason in ways humans could not possibly reason—it just does not claim or attempt to replicate human legal reasoning. Indeed, there is at least one way in which my system's reasoning is inconsistent with evidence regarding human commonsense and everyday reasoning: there is strong evidence that humans generally use higher-order relational information to learn concepts, reason about the world, and construct arguments, but my system learns and reasons from relational cases that lack those higher-order structures connecting case facts to outcomes. It is not clear that this is inconsistent with how humans perform *legal* reasoning, since legal reasoning is a specialized form of human reasoning that does not necessarily track with other forms of human reasoning, but the cases upon which the system is tested lack the higher-order structures that are known to be useful in non-legal domains, and the system therefore does not make use of them. Perhaps future evidence about human legal reasoning will eventually reveal that legal reasoning tracks closely with other kinds of reasoning, and that therefore the mechanisms by which the present system operates do not align with how lawyers discern and apply legal rules. Or perhaps future empirical evidence will show the opposite, and demonstrate that the engineering choices made in service of building a functional task-oriented system happen to track with how legal reasoning works in humans.

This thesis is agnostic on future evidence regarding human legal reasoning. Still, it is important to describe five fundamental premises I held that informed the design of the model described in the subsequent chapters. These are: lack of evidence regarding human legal reasoning; a desire to demonstrate that the system is actually learning legal rules from cases; the unreliability of legal case explanations; the theoretical scalability of the system being proposed; and limitations on CNLU, the natural language understanding system used to generate the dataset. I address each in turn.

Because the nature of legal reasoning is hotly debated, it is unclear what exactly is to be modeled if one seeks to build a faithful model of human legal reasoning. That is, what legal reasoning actually *is* is still a matter of vigorous disagreement. For example, while everyone agrees that as an empirical matter lawyers in their briefs and judges in their opinions use analogies to past cases, there is strong disagreement about whether those analogies play a role in helping the jurists resolve the outcome of the case, or whether they are purely rhetorical devices meant to illustrate the rule being applied and stave off any claims of arbitrariness (because decision-making consistent with past cases reduces judicial discretion). There is little evidence and no consensus view on how legal reasoning proceeds *in situ*. Much of the analysis has instead focused on legal writing, which is arguably not legal reasoning at all but rather the product of a legal reasoning process that does not necessarily track with the reasoning itself. Because there is no consensus view on the nature of legal reasoning, I do not attempt to implement any given model of it. Doing so is not necessary to present a pure AI & Law system that accomplishes the task of extracting and applying precedential rules.

Additionally, if legal reasoning in humans does depend on the presence of higher-order structures, the inclusion of those structures in the training cases would undermine the claims of the thesis by providing the system with that which it is meant to learn. That is, the structures in question are the justifications tying legal outcomes to legal case facts, and as such are exactly the explanations that the system seeks to learn; providing them to the system undermines the claim that it is learning them. The legal rule announced in a common law case can be seen as nothing more than the facts of the case connected with its legal conclusions. "Because Bob was on Alex's property without permission, Bob trespassed" is both a connection of facts to conclusion and just about the most concise statement of the trespass rule possible, even if it is instantiated with the entities from a given case. The connections between facts and conclusion—the explanations—*are* the rules, at least in the sense that they are that which becomes binding on future cases. (Certainly those explanations are not rules in the formal logic sense of the word “rule”.) So were the explanations included in the cases provided to the system and from which it is supposed to learn, the system would be getting fed exactly the information which it purports to learn, which would undermine the claim that it was learning it. Including those explanations would allow me only to have shown that the system is able to abstract facts explicitly given to it. If the explanations are not provided, then I can reasonably claim that the system is discovering those explanations—those rules—for itself across cases. The first finding would only demonstrate that an abstraction system works, which has been demonstrated repeatedly over more than a decade of research on the SAGE system; the second makes a strong claim about *discerning* legal rules, not just abstracting instantiations of them.

Additionally, although human judges and lawyers may well rely on the explanations provided in in common law cases, those explanations are nonetheless inherently suspect. Whoever explicitly associates case facts and conclusions is not always assumed to have done so properly by later judges and advocates. For example, the law of trespass has sometimes been held to require that the property trespassed upon was damaged: a court would say, "Because Bob walked across Alex's property without permission, and because in doing so Bob destroyed Alex's begonias, Bob trespassed on Alex's property." But courts eventually stopped requiring that the property be damaged, and when they did, they did not claim they were changing the rule, but that the prior courts had not properly understood what rule they had been applying, because the prior court did not accurately identify which facts in the case were relevant to its conclusion. (Note the connection to the previous point about rules being no more than facts bound to conclusions.) The new court would say something like, "Here, Dan walked across Carl's lawn without permission. Although the prior court in *Alex v. Bob* said that Bob's damaging Alex's begonias was relevant to Bob's trespassing, it is clear that the mere fact of Bob being on Alex's property without permission was sufficient to establish a trespass. Therefore, even though Dan did not damage Carl's property, he committed a trespass." Thus the rules announced in a case are not necessarily the rules that are later understood to have governed that very case. Learning from common-law cases that only include case facts and case outcomes (not explanations) makes a stronger claim that the learning system discerns the true legal rules in a common law domain, rather than what the judge in any given case said they are. Note that this claim is weakened when dealing with statutory interpretation, where the rule written in the statute is required to be given meaning, although even statutory interpretation is subject to common-law interpretation. Nonetheless, because the present research was done in a pure common law domain—specifically to avoid dealing with statutory interpretation—the objection is irrelevant to the current work.

My engineering commitments were also informed by practical concerns regarding what comes next for my system. To ever be useful to people in the real world, the system must be able to scale beyond the cases I have provided to it. I therefore sought to establish a workflow that contains only steps that in theory could be performed by crowd-workers. For any given case, this includes (1) identifying the case facts and outcomes, (2) simplifying them to the syntax that can be handled by CNLU, (3) selecting amongst CNLU choice sets interpreting words and phrases, and (4) resolving common reference tokens[21] to each other. (For my thesis I have also had to supplement CNLU's linguistic coverage and write narrative functions; I believe these functions could also be crowdsourced, but there will also come a point where CNLU's linguistic and narrative function coverage is sufficient to make having crowd-workers do so unnecessary; hopefully steps (3-4) can eventually be obviated as well). Identifying case facts is a much simpler task—and is therefore more able to be converted into a task for crowd-workers—than identifying the justifications for the outcomes given the facts. Case facts consistently appear at the beginning of a judicial opinion (after jurisdictional and procedural statements) and are therefore easily identified, whereas the statement of the rule can appear almost anywhere in the much bulkier and

[21] For example understanding that, given the sentences "Isabel picked up the rock. It was heavy, and she grunted," the pronoun "it" refers to the rock while "she" refers to Isabel.

less consistently structured analysis section. Sometimes the explanations in an opinion are underspecified; sometimes they are over-specified. Furthermore, identifying the statement of the rule is exactly the job of a lawyer, who has specialized training in doing so. Because this is a specialized skill, crowd-workers should not be assumed equipped to isolate a judge's explanations from her opinions and ensure they are properly translated into the system's internal representation system. Therefore, in service of the future scalability and impact of this system, I made sure the system could learn without the explanations that I was not confident would be available to it in future cases.

This scalability point is connected to a final note about the nature of the dataset. While this point is not a substantive reason that would apply to any theoretical legal learning system, it was an important practical consideration: CNLU is not currently capable of generating the higher-order facts connecting case facts to case outcomes. These explanations will generally include multiple subordinate clauses, each of which refers to a different entity or event from prior sentences. (For example: "The plaintiffs did not trespass on the track because the defendants' allowing public access to the track and their knowledge that the public used the track as a road implied permission for the plaintiffs to walk on the track.") If those representations were required then they would need to be hand-generated, which is a difficult task for crowd-workers and would undermine the fourth point above. It also would undermine a claim of the thesis, which is that my system is the first legal case-based reasoning system that uses logical representations of cases that also does not use extensive hand-encoded representations. I have drawn a bright-line rule against writing any case fact logical expressions by hand, and I have not crossed that bright line.[22]

[22] Case *conclusion* statements were hand-generated because they too require referring back to several different entities across many sentences in the case, but these (a) are not statements of case *facts* from which the system is

I am a firm believer in the usefulness of cognitive modeling. I believe that evidence regarding how humans engage in legal reasoning would be extremely useful, and one possible continuation of this thesis work would be to modify my system to be in line with such evidence. But the question "Can a legal reasoning system be engineered that is both faithful to human legal reasoning and accomplishes a legal decision-making task?" is a big research question. My thesis instead asks a smaller question, focused only on accomplishing the task itself, without regard to whether its operations accord with evidence on human reasoning.

C. LEGAL ACADEMIC CONTRIBUTION: A NEW PERSPECTIVE ON ANALOGY

As described in Chapter 2, legal academics have extensively debated the role of analogy in legal reasoning. But perhaps surprisingly, as far as I can tell the analysis has always and only concerned analogies between some case at bar and one or more individual prior cases. That's surprising because evidence from Psychology has long shown that analogical reasoning is not only useful for reasoning about some new situation, but for learning across similar situations. Therefore, while I do not claim that my thesis research provides direct evidence that lawyers and jurists discern legal rules by comparing precedent cases to each other, this thesis nonetheless provides the opportunity to propose a heretofore unexplored role for analogical reasoning in legal precedential reasoning.

I hypothesize that, regardless of whether analogy helps resolve an individual case with reference to a precedent, analogy can play a role in determining what rule governs a whole line of cases. Here is the fundamental insight behind the hypothesis: when lawyers or judges are looking for a precedential case that might govern some case at bar, they are likely to examine several cases

learning its legal principles, and (b) are single non-nested statements that always follow the same format—to the point that they could be generated using drop-down menus—and could therefore be generated by crowd-workers.

in a given legal doctrine looking for the “right” precedential case—the case most similar to the one being reasoned about. And after they examine each precedent case and move on to the next one, they do not magically wipe their memories and forget the case that they have just moved on from. That is, it is natural to assume that, as lawyers consider a new precedent case in light of the case at bar (to see whether it is the appropriate case to draw analogies from in the opinion or arguments), they also consider that precedent case in light of all the other precedent cases they have already seen. And if so, those lawyers would have trouble *not* building up a model of the rules governing that series of cases, because generating schemas of like things is something that comes naturally to humans and is indeed a fundamental mechanism of human concept learning (Gentner & Smith, 2013).

I hypothesize (and my research system is built on the assumption) that analogical comparison and generalization naturally extracts legal principles from precedent cases, and those principles are then converted into rules. Note that this is related to the middle-ground position staked out by Levi, Sunstein, and Brewer, but it fills in a critical piece missing from their accounts: how the rules governing the past cases are themselves discovered. But it is also not inconsistent with Alexander & Sherwin’s criticism of the idea that analogies to past cases are used to resolve new ones. Indeed, Alexander & Sherwin may allow for a larger role of analogy in legal thought than they initially appear. For one thing, they admit that a concept of relevant similarity is crucial in determining whether like cases should be settled alike, they just do not characterize that determination of similarity as being analogical reasoning. But from a cognitive scientist’s perspective, the fact that a lot of the work in constructing a good analogy is in determining what should be aligned with what—in determining the similarity of two cases—does not mean that the reasoning supported by such a determination is not analogical reasoning. Furthermore, Alexander and Sherwin focus on analogy exclusively as a form of case-to-case reasoning. They say that "*searching* for analogies and common principles that link past and present cases . . . might play a useful role in the development of common law." That is, they admit a role for analogy in aiding a legal reasoner to discover rules. I would go one step further, and argue that is exactly the process through which jurists search for analogs to present cases that yields the common principles across the past cases, and that they are yielded through the successive constructions of analogies and analogical generalizations across those past cases. Thus it is possible that the disagreement is largely definitional: I use "analogy" in the sense in which it is commonly used in Cognitive Science, to refer to the often intuitive and cognitively automatic alignment of like scenarios, rather than the more laborious and explicit construction of analogies of which lawyers often speak.

Before turning to the model I developed and the experiments I performed on it, I will briefly describe some properties that an automated legal reasoner ought to have, and why.

D. DESIDERATA OF AN AUTOMATED LEGAL REASONING SYSTEM

Ideally a legal reasoning system would have several qualities to support not only its functionality but its usefulness. Many of these ideas build upon Ashley's proposed requirements of an ontology for legal case-based reasoning (Ashley K. D., 2009). Ashley proposed that, among other things, a legal CBR ontology must include a hierarchy of factors to represent case facts at a variety of levels of detail and abstraction. I agree with Ashley, and offer some observations on some subjects where I believe he does not go far enough.

i. RELATIONAL REPRESENTATIONS

The first feature an automated legal reasoning system ought to have is a representational system that supports expressing and understanding relationships between the entities and events in the case. This representational system might be first- or higher-order logic, but the key is a to go beyond lists of features. That is, an automated legal reasoning system should have representational capacities beyond feature vectors of the type that feature prominently in machine learning research.[23] To represent much of the important, salient information in legal analysis as features (or "factors", to use the word Professor Ashley has proposed for legal features) involves applying sophisticated judgment to a complex set of relational facts, judgment that may not be easily automatable. For example, in analyzing a Tort claim of assault, one might care a lot about the factor of one person threatening another person. And certainly that could be represented as a feature: DefendantIsThreateningPlaintiff can either be a 0 or a 1. But that feature itself encodes a substantial amount of structure: the feature describing the situation is dependent on subtle differences in the structure of the facts describing the situation. We might say that the feature is present, for example, when a 6'4" person takes a newspaper and tightly rolls it up, grips it in by one end, and stands very close to a 5'2" person, huffing and puffing onto them. But if almost the exact same thing happens on a crowded subway train at rush hour, the feature probably is not present. Why? Because of the relationship of the actors and actions to the background context. Even if these can be reduced to a feature, it is not clear how to do so in a consistent and automatable way, nor how to preserve the flexibility needed to differentiate two highly similar cases like the ones just described, without using relational representations.

The work on HYPO and its progeny, and all the factor-based research, has depended on humans providing those features, and on those humans doing the work of reducing the actual structured, relational information down to the features themselves. (Ashley & Bruninghaus (2006) automatically tagged cases with relevant factors, but the factors themselves were still identified by

[23] I do not wish to imply that Ashley, given his work on factors, is hostile to this point: in his 2009 paper, he notes that relations describing case facts or the relationships between the facts and outcomes might need to be represented in a structured way, such as a semantic network, as Branting did in GREBE (Branting L. K., 2003).

hand, just attached automatically.) For a statute like "no vehicles in the park," perhaps a legal reasoner could get most of the way to a decent comprehension of the law through pure features. And perhaps reducing structure to factors might always provide a path forward for a factor-based system. But if the factors cannot be specified beforehand, or if the factors might not be assumed to be perfectly understood (or to shift over time) then factors will be unable to capture the complexity of information in legal cases (McCarty, 1997), and representing situations with all the complexity of the structures and relationships involved will be critical. Regardless, reducing the structures to features involves moving the goalposts, because one still must determine when a situation involves a particular feature.

ii. Explainability and Transparency

Reducing structures to features also impacts the desiderata of explainability and transparency. Understanding why a legal case came out the way it did is a critical component of jurisprudence and the rule of law.[24] Legal arguments and decisions made by humans need to be understandable so the that arguments can be comprehensible, so parties can understand why their case was settled as it was, so a reviewing court can examine and review the lower court’s decision, and so a future court can understand and apply the case as a precedent when it is on point for the future case. Even if automated legal reasoning systems are not used for legal decision-making or to replace lawyers outright, to the extent they will be useful in a supporting role in the legal system they must have these properties as well.

The “explainability” of a legal argument might mean two different things: the comprehensibility of the argument itself, and the explanation as to why the advocate is making the

[24] As Pasquale (2015) points out, transparency is not necessarily an end in itself, but as a step towards interpretability, which is the ultimate goal.

particular argument that it is. The need for explainability in legal argument under the former meaning is almost tautological. What is an argument except an explanation as to why a particular outcome should be derived? It is hard to imagine how an argument could be unexplainable—meaning incomprehensible—and still qualify as an argument. But the latter meaning is a little trickier. For example, an argument might be made pretextually: the system could derive outcome X using a rule A→X, but understand that B→X arguments tend to be more successful, and so present argument B for outcome X even though it believes A is a better argument. Similarly, an argument in favor of an extreme and undesired outcome might be used as a rhetorical strategy designed to result in some unstated middle ground. Should an automated legal reasoner that generates arguments have the process by which the arguments are generated be traceable and inspectable?

Arguments are possible on both sides. The strongest argument against, I believe, is that the argument itself should stand on its own regardless of why that argument was generated. That is, if an argument is consistent, coherent, and persuasive, what does it matter if the system that generated it did so pretextually, or in a manner that we would call cynical if a human had done it? As long as the first meaning of "explainability" is met for some argument, this perspective goes, then the second is irrelevant. So if some Large Language Model can fluently generate a persuasive paragraph in favor of some outcome, there is no need to be able to crack open the internals and verify that the system actually "believes" that which it is arguing for. After all, we do not require human lawyers to believe that the arguments they are proposing on behalf of their clients are the best arguments concerning the entire situation. And sometimes pretext is openly accepted within our legal system: pretextually holding a suspected murderer overnight on a traffic violation, for example, or sending Al Capone to jail for tax evasion.

The argument in favor of having automated systems' arguments be traceable to their sources is that computers are different from humans, and we should make sure they are not lying to us or generating random nonsense before adopting their arguments. And with computers, unlike with humans, we have the capacity to insist that their reasoning be inspectable and traceable before we are willing to entertain it as persuasive. Given the extensive evidence of bias in machine learning systems trained on real-world datasets, it is safer to assume that an argument from such a system is biased than that it isn't. While we might have no option but to accept that human lawyers will sometimes make pretextual arguments, pretextual arguments are generally bad from a moral perspective, and we should not blithely propagate them within our legal system.

I believe the argument in favor of requiring explainable arguments is the stronger one. If automated legal reasoning systems are going to be generating arguments, they should do so in such a way that human users can understand *why* a given argument was generated, so they can decide whether they want to open themselves up to being persuaded by the argument or to promulgate it in a legal proceeding. This is true even if the argument falls within one of the acceptable forms of pretext: If some system generates the argument "hold him overnight for a traffic violation" because it knows it will not succeed with the argument "hold him overnight as a murder suspect," it is still useful for the human overseeing the system to understand that the system views the detainee as a murder suspect. It is one thing to tolerate an AI system manipulating a situation in ways that are legally acceptable, and it is another for the AI system to manipulate the people overseeing and managing that situation. And explainability becomes even more important when moving from legal argument to legal reasoning and decision-making.[25]

[25] The explainability of a system depends on who is doing the interpretation. Generating natural-language descriptions of a system's predicate-logic internal representations is an unsolved problem, so a system that

The rule of law means that law is not applied in an arbitrary or ends-oriented way. Procedural justice depends on the perception of the fairness of some legal process. Both of these in turn depend on being able to understand not only *what* happened in some case, but *why*. Of course, those are properties required by actual legal decision-makers, and as argued previously, legal decision-making should not be delegated to computers. But the property of explainability is desirable even in legal systems that only support decision-making. Most importantly, one must ask, what is the role of supporting decision-making? It can only be to help someone come to their own best decision or understanding about a case: to do more would be to actually be doing the decision-making that is meant for the human, and to do less would be not to support the decision-making process. That is, if the system is not generating an explanation for the decision it recommends, then the human receiving the recommendation has the choices only to either reject the recommendation, or to accept it and thereby substitute the system's judgment for their own. And if the system *is* generating an explanation, then the explanation should be detailed and accurate for the same reason that explanations of legal arguments should be. The point is not only to be persuasive but to be *correct*, and the correctness depends on inspecting the explanation. That's true regardless of whether the human in question is an actual legal decision-maker trying to understand a case or the law governing it, or a potential litigant trying to understand the legal landscape in which she is considering taking a step.

This holds true when a case is considered in the larger context of the legal landscape, that is, beyond the impact of the case itself on the litigants. For a higher court to review a lower court's

generates explanations in predicate logic is not truly explainable without a human domain expert who can interpret those representations. Still, a system that can generate a reliably accurate explanation that can be interpreted by *some* human is explainable in a way that systems – like modern large-language models – that generate potentially inaccurate explanations, or ones that no human can reliably understand, are not.

decision the lower court does not necessarily need to have provided an explanation, but one would be helpful. Without the explanation the higher court must reconstruct for itself what the lower court might have been thinking, and may do so incorrectly, or in a way that is uncharitable to the lower court. If the lower court is to protect its decisions from reversal, it should explain why it arrived at them.

Similarly, without explanations as to why the facts led to an outcome, the rule of a case will be confined to its facts and its outcome. This can already be quite useful: indeed, the cases from which the system in this thesis learns are no more than facts and outcomes! Treating cases as only facts and outcomes allows future courts to reinterpret what it was that the past court did in a given case. But providing the explanations allows a case to be extended beyond its facts and can instruct lower courts how to proceed in novel cases that are not directly constrained by the case being settled.

iii. Scalability

The last desired feature of an automated legal reasoner that I will discuss here is scalability.[26] For an automated reasoner to be useful it must be applicable in more than a small number of situations. It might well be domain-specific, but it should be applicable to the largest possible set of cases within that domain. If it is domain-general, it should be able to be expanded to new domains in a relatively straightforward way. Both these goals require the system to be scalable and to be able to be modified after its original launch.

[26] O'Neil (2016) lists scalability as one of the three criteria that define what she calls a weapon of math destruction, an algorithm with the potential to cause great societal damage (the other two are opacity and the capacity to cause damage to the subjects of the algorithm). It is my hope that these systems being transparent and helpful to their users and subjects would prevent them from being harmful, but I acknowledge that the risk exists.

Automated legal reasoning systems will either have their legal reasoning mechanisms set in advance by programmers, or will learn from a dataset, or both.[27] The law changes as a result of legislative intervention, judges changing the rules for known situations, or new situations presenting themselves. If the reasoners are predefined by their programmers, then there must be a mechanism by which programmers can intervene when the law changes, or write rules to deal with new kinds of situations. If the reasoners learn from a dataset, then that dataset must be able to be updated as new cases that change the rules come in. (Ideally there would be a combination of both, to allow a system to both track changing case law and to recognize when a particular event has changed a broad swath of the law at once.) Either way, for a legal reasoning system to stay current and useful, there must be a mechanism by which new legal information can be brought into the system.

The desired feature of *scalability* refers to the ability of engineers or users of some legal system to expand its functionality over a greater range of information or domains. What is required will look different system-by-system. Perhaps it will refer to the ability to add more cases to a training dataset, either to more clearly delineate the boundaries within some legal domain, or to expand the system to multiple legal domains. Perhaps it will refer to the ability to have new rules input into the system and tested to ensure their consistency with the existing case law and doctrine. Or perhaps it will refer to little more than adding new forms for users of a system like DoNotPay to be able to contest parking tickets in a larger number of bureaucratic systems. Regardless, if a legal system accomplishes some task, it should be engineered with the goal of having every possible user of the system accomplish every possible permutation of the task (or else it will

[27] There are reasoning systems that can learn rules from a stream of inputs rather than from a predetermined dataset; I count these as “learning from a dataset,” even if the data are not fully specified at implementation time.

unfairly favor the subset that happen to be supported). Because it is unreasonable to expect any system to launch with that extent of domain coverage, system engineers should ensure that their systems are scalable and can be supplemented after launch.

I now turn to the research itself that I accomplished for this Thesis, beginning with the Illinois Intentional Tort Qualitative Dataset.

CHAPTER 4: THE ILLINOIS INTENTIONAL TORT DATASET

In Common Law legal systems, legal cases are resolved not only with reference to statutory rules, but also to prior cases concerning similar legal claims. While precedential reasoning is not all of legal reasoning and jurisprudence, much research in AI & Law has sought to formalize and model common law legal reasoning. In turn, researchers developing such systems have used real-life common-law cases both to inspire their models and to test them. Datasets of such cases, and especially the formal commitments and assumptions made by those who collected those datasets, have therefore been important to the AI & Law research community. This chapter presents the first publicly-available dataset of historical tort cases represented in predicate logic extracted directly from the judicial opinions' statements of facts using a natural language understanding system.[28]

Statistical machine learning techniques have led to the development of large-scale databases collecting cases in their original linguistic form. Caselaw4 (Petrova, Armour, & Lukasiewicz, 2020) annotated 250,000 cases with information about their courts, decision year, which legal doctrines were implicated, etc. The Vaccine Database was used in developing LUIMA (Grabmair, et al., 2015), a retrieval and reasoning system that uses IBM Watson to reason about

[28] The dataset is available as of 2023 at https://www.qrg.northwestern.edu/Resources/caselawcorpus.html.

vaccine injury cases. But while researchers using large-scale machine-learning techniques have recently developed legal retrieval and reasoning systems that operate over raw text (Branting, et al., 2019), AI legal reasoning systems have often required formal machine-interpretable representations beyond the texts furnished by courts, and researchers have had to annotate their cases with those representations.

AI & Law's long history of using symbolic case representations demonstrates how useful such representations are to computational legal models. Symbolic representations are a natural fit for legal reasoning, because the rationality, explainability, consistency, and transparency of reasoning that symbolic representations support are considered hallmarks of good *human* legal reasoning. But in an era of big data and large language models, research that relies on human annotation is unlikely to gain much traction in the real world or to be perceived as advancing the state of the field. Semantic interpreters – systems that take in text and output symbolic representations thereof – may provide a solution. These systems can split the difference, operating over natural language directly (and thus at scales beyond hand-encoding), but providing the rich representations researchers need to create subtle and interpretable models of legal reasoning. The Companion Natural Language Understanding system (CNLU) (Tomai & Forbus, 2009) is such a semantic interpreter built into the Companions cognitive architecture (Forbus & Hinrichs, 2017), and is described in Research Background.

A. THE ILLINOIS INTENTIONAL TORT QUALITATIVE DATASET

The state of Illinois was chosen for the dataset's jurisdiction because it is the graduate institution's home state. CNLU is not currently able to handle the linguistic complexity and legalistic formalism of statutes,[29] so I sought pure common-law doctrines whose rules are expressed in plain English (however complex) in judicial opinions. The doctrines identified were in Tort: Trespass, Assault, and Battery, and the affirmative defense of Self-Defense.[30]

Cases were collected by searching WestLaw and LexisNexis for the specific doctrines in Illinois from the late 19th Century onward. Retrieved cases were traced forward and backwards, i.e., the precedents upon which those cases depended, and the subsequent cases depending upon them, were retrieved. Cases overturned on appeal, unrelated to prior or subsequent cases, in esoteric and rare legal areas, or with unreasoned decisions, were excluded. Cases were limited to those occurring after 1870, because before that date decisions were often structured informally, with facts less clearly defined and offset from conclusions.

Collected cases were organized according to the doctrines they illustrate and annotated with their case reporter index, decision year, and court. Case facts and conclusions were manually identified and stored as a string argument of a fact tying them to the case. Eleven of the case opinions were from appellate courts resolving an issue of law for the lower court, where the appellate court laid out alternate set of facts left ambiguous by the lower court, along with the associated possible conclusions. These cases were converted into two cases for the dataset, one for each alternative set of facts and corresponding conclusion. The dataset thus comprises 88 cases illustrating 112 distinct legal claims in tort, including both positive and negative examples of claims (a positive example being one where the court found that the claimed tort or defense had in

[29] That said, researchers have studied automatically constructing semantic interpretations of statutes (Wyner, Gough, Levy, Lynch, & Nazarenko, 2017).

[30] The most common Tort claim is Negligence, which was excluded because it turns on the squishy open-textured concept of the "reasonable person." Tort law professors often define "reasonableness" as "whatever the jury thinks is reasonable." Negligence therefore presents one of the most difficult cases for an automated reasoner. Because intentional torts are substantially more straightforward, this initial dataset focuses only on them.

fact occurred, and a negative example being one where the legal standard was not met). These include 17 assault cases (12 positive, 5 negative), 40 battery cases (30 positive, 10 negative), 43 trespass cases (29 positive, 14 negative), and 12 self-defense cases (7 negative, 5 positive). Positive cases outnumber negative ones because positive cases are more likely to be published and later relied upon as authority.

CNLU is currently unable to handle the complexity of judge's descriptions of case facts, which often include run-on sentences punctuated by series of semi-colons, comma-separated lists, long descriptions, and asides—as this sentence does—so the case texts were simplified such that CNLU could understand them (Figure 4.1). The simplification process worked as follows: first, the parties' names were reduced to their party designations ("the plaintiff(s)/defendant(s)"). Specific identifying information about locations (e.g., street names), parties (e.g., names of their business, their spouses, etc.), prices paid for things, and dates, were removed. For cases with multiple causes of action, facts identified as only being relevant to a claim other than the tort in question were removed. Words not in CNLU's vocabulary were replaced with synonyms or added to the vocabulary. Finally, longer sentences were broken down into their simplest clauses; complex grammatical structures were rephrased to use simpler constructions; and compound nouns were

Figure 4.1: Original vs. Simplified Case Facts. Bishop v. Ellsworth, 91 Ill. App. 2d 386 (1986)

Original Text, Bishop v. Ellsworth: "On July 21, 1965, defendants, Mark and Jeff Ellsworth and David Gibson, three small boys, entered [the plaintiff Dwayne Bishop's] salvage yard premises at 427 Mulberry Street in Canton, without his permission, and while there happened upon a bottle partially embedded in the loose earth on top of a landfill, wherein they discovered the sum of $ 12,590 in United States currency. [The] boys delivered the money to the municipal chief of police who deposited it with [the] Canton State Bank. The defendants caused preliminary notices to be given as required by Ill Rev Stats, chapter 50, subsections 27 and 28, (1965)."

Simplified Text: "The plaintiff owns a salvage yard. The defendants are young boys. The defendants entered the salvage yard. The plaintiff did not permit them to enter the salvage yard. The defendants found a bottle containing $12590 of money in the plaintiff's salvage yard. The defendants brought the money to the chief of police. The defendants placed notices about the money in the newspaper."

Figure 4.2: "The plaintiff climbed to the balcony" in Neo-Davidsonian Event Representation.

```
(isa balcony4180 BalconyLevelInAConstruction)
(isa climb4579 Climbing)
(isa plaintiff4574 Plaintiff)
(performedBy climb4579 plaintiff4574)
(to-Generic climb4579 balcony4180)
(isa plaintiff4574 Plaintiff)
(toLocation climb4579 balcony4180)
```

sometimes rephrased as declarative sentences (e.g., a long sentence referring to "the defendant's salvage yard" became several sentences including "The defendant owns a salvage yard."). After simplification, texts were passed into CNLU for semantic interpretation. To ensure maximum fidelity to the text, choice sets were selected manually rather than automatically.

Legal reasoning operates over complex real-world situations, so a rich, accurate understanding of legal texts is critical. Narrative functions can be used to make inferences about the meaning of a sentence that is not captured in the strict semantics of the words. To illustrate: given the sentence "the plaintiff climbed to the balcony," CNLU might yield the representations shown in Figure 4.2: a climbing event, done by the plaintiff, with its endpoint at the balcony (indicated by the toLocation statement). But missing from this description is that *the plaintiff is now on the balcony*. Trespass in its simplest form involves someone being on private property without permission, so this missing fact is critical to understanding that the plaintiff might be trespassing in this situation. Similarly, it is impossible to understand anything about Assault or Battery without understanding what actions constitute threats or physical contact. But this information is so obvious to humans that it is rarely explicitly stated in an opinion's statement of fact. Judges rarely explain that a criminal pointing a gun at a policeman are threatening the policeman.

Narrative functions can make these kinds of commonsense inferences at the level of language processing and understanding. The fact that narrative functions operate within CNLU is critical: they are part of language understanding, not logical rules applied to its output. Narrative functions written for some task can be reused by CNLU system for other domains, if appropriate. Because legal reasoning operates over real-world situations in all their complexity, a rich and accurate understanding of the text describing those situations is an important aspect of a legal dataset generated from a bare description of events.

I wrote narrative functions to make commonsense inferences in a variety of situations that frequently recurred in the dataset's cases (approximately two hundred narrative function detection rules for 93 narrative functions). The bulk of these were rules to infer (1) where objects were (as in the climbing example), (2) whether an event causes damage (e.g., that stabbing something damages it), (3) transitive ownership (e.g., if you own a building and the building has a balcony, you own the balcony), (4) whether an event involves two things touching each other, (5) part/whole relationships, and (6) that certain actions create new entities (for example, Adam suing Bob creates a lawsuit entity, with Adam as plaintiff and Bob as defendant).

To ensure I was writing commonsense language understanding rules and not legal rules, I avoided writing rules that are themselves legal conclusions. The case opinions helped determine what is a legal versus a commonsense conclusion. If an opinion indicated something was true as a matter of law, it was off-limits. I also asked whether a rule would apply in circumstances other than a legal proceeding, and only wrote rules for inferences that would. Nonetheless, determining what qualifies as a commonsense inference versus a legal inference is tricky and can be seen as a limitation of my approach, discussed below.

B. NOTES ON THE DATASET

The first five cases in the dataset are included in Appendix A; the entire dataset is not appended because doing so would make this document over 1000 pages long. An evaluation of my experimental reasoning techniques using the dataset is presented in Chapter 6: Here I offer some observations about the dataset drawn from those analyses. I hope the reader will indulge my reference to lessons drawn from results before those results themselves are presented.

The greater number of positive cases than negative cases in the dataset may bias statistical methods towards simply accusing the defendant of being guilty and moving on, and generally performing well. (This lopsidedness simply reflects the fact that fewer opinions finding someone not liable of trespass are published, particularly at the appellate level, where most published opinions are found). This risk is illustrated by the fact that the statistical baseline methods outperformed one of my approaches (Analogical Reasoning with Positive Generalizations, or ARPG) in positive cases but not negative cases. A similar risk comes from the fact that, in most cases, it is the defendant that is accused of tortious behavior. This is not true in all cases: for example, a significant number of trespass cases in the dataset raise the issue of trespass as a defense against liability, where the plaintiff was injured on the defendant's property, and the defendant seeks to avoid liability by accusing the plaintiff of trespassing. To assess whether a reasoning technique actually understands what it is doing, special attention may therefore be paid to negative cases and cases where the parties do not stand in the standard relation to each other.

The limitations of the language system and of the process by which commonsense narrative function rules were generated must be acknowledged. The predicate logic representations of the dataset do not yet reach the goal of being generated by feeding raw legal text into a language understanding system, because to my knowledge no language understanding system exists that can reliably both handle the complexity of legal text and generate accurate symbolic logical representations from them. As it relates to the goal of being able to generate datasets automatically from legal text, CNLU features three limitations, each of which are areas of active research in the Qualitative Reasoning Group as well as others.

First is the syntactic complexity of the language handled: CNLU still relies on a human taking the original text and simplifying. The complexity of surface forms CNLU can handle has progressed over the years, but while the goal is that it eventually can handle arbitrarily complex grammatical English input, that goal will not be achieved this year or next. In the meantime, one possible solution is that a large language model might be trained to simplify texts to a level CNLU can understand, eliminating the need for manual simplification.

Second, to create this dataset the choices from choice sets generated by CNLU were selected by hand to ensure maximum semantic fidelity to the original text. Much work has been done and continues to be done in the Qualitative Reasoning Group to enable CNLU to automatically select the appropriate choice sets (Barbella & Forbus, 2015; Ribeiro & Forbus, 2021). I made the choice to hand-select the choice sets to ensure that the generated representations were as faithful to the original meaning as possible, because the purpose of creating the dataset was not to evaluate the language system but to create an accurate legal dataset for use in AI & Law research. That said, the choice-set selection system already works sufficiently well that I could have used it at scale were I willing to take a trade-off in semantic fidelity.

The third limitation is the most severe one, and it is the need for something like narrative function to express not only what the words of a text literally *say*, but what those words *mean*. Real-world flexible common-sense reasoning is one of the most persistent and intractable problems in AI research, and I do not believe it will be solved in the general case absent a significant leap forward in AI. In the meantime, the options are to try to create generally-applicable rules that the language system can use across domains and applications, or to simply accept that facts that are obvious to humans will simply remain unknown to the computer. I believe that, in the domain of legal reasoning where such facts are critical to understanding what has happened in a case and why its judicial outcome therefore follows, accepting that those facts remain unknown to the system means guaranteeing that a computer system will either be unable to learn legal concepts, or will learn the wrong ones. I invite disagreement and discussion on this point.

C. FUTURE WORK

There immediate area of future work relating to the dataset itself is to expand the dataset. During the collection and simplification stage of creating the dataset, I collected over one hundred cases in Illinois Contract Law illustrating issues of contract formation, breach of contract, and the defenses of impossibility and duress. These cases are already simplified and ready to be processed by CNLU. They are missing only due to the pressures of time relating to my graduate studies. The addition of contract cases would allow the Illinois Intentional Tort Qualitative Dataset to become the Illinois Common Law Qualitative Dataset and open up new avenues of AI & Law research.

CHAPTER 5: CONCLUSION-VERIFIED ANALOGICAL SCHEMA INDUCTION.

For reasons that will be discussed in this chapter, the cases in the Dataset appear sufficiently dissimilar to each other according to the Structure Mapping Engine's scoring systems that relying on using SME's similarity score to control schema assimilation in SAGE was ineffective. More precisely, the SME scores were essentially the same regardless of whether the cases were ones I thought should generalize together or not. I developed a new technique, CASI, to determine when to assimilate cases into a generalization. This chapter describes the technique and situations in which it should prove useful. It also describes the generalizations built using CASI over cases in the Dataset, generalizations used by the rest of the experimental systems described in the thesis.

A. INTRODUCTION AND BACKGROUND

Learning useful concepts from examples is a core area of research in Artificial Intelligence. This reflects the fact that the ability to derive concepts from individual experience is a hallmark of human intelligence. AI researchers have undertaken an enormous variety of approaches to this task: demonstrating how to solve tasks to robotic learners (Chernova & Thomaz, 2014), learning logical definitions from instantiated relational examples (Quinlan, 1990), learning tasks from observations of human performance (Gulwani, et al., 2015), learning hierarchies of concepts using Bayesian probabilistic reasoning (Grant, Peterson, & Griffiths, 2019), comparing positive cases to near-misses to strengthen category boundaries (McLure, Friedman, & Forbus, 2015), (Rabold, Siebers, & Schmid, 2022), learning event schemas from graphical event representations using neural networks (Jin, Li, & Ji, 2022), and many more. And of course, many of the recent advances underlying Deep Learning can be understood as using examples to learn underlying concepts (LeCun, Bengio, & Hinton, 2015).

In some cases researchers might have few preconceived opinions about the concepts underlying their data and the purposes for which those concepts can be used – they toss their data into a machine learning algorithm and see what concepts the algorithms reveal (Bengio, 2012). In other cases the researchers have specific concepts in mind, and use supervised learning to learn what underlies those concepts (Krizhevsky, Sutskever, & Hinton, 2017). This is particularly useful when developing systems designed to reason about and solve problems: researchers might know that particular concepts are core to a problem domain, but rely on their learning algorithm to carve the boundaries of those concepts for generalized reasoning purposes (Fitzgerald, Short, Goel, & Thomaz, 2019).

This chapter presents a modification of the Sequential Analogical Generalization Engine (SAGE) (Kandaswamy & Forbus, 2012), a system that uses analogical reasoning to generate schemas from examples that share underlying structures. SAGE's design is inspired by psychological evidence regarding how humans learn concepts through comparison. SAGE learns concepts by positive examples, in the form of structured, relational examples, added incrementally. As usual with structure-mapping (Forbus, Ferguson, Lovett, & Gentner, 2017), higher-order relations help indicate which lower-order relations are relevant to conclusions, and the mapping of deeper structure drives the quality of the match. But what about when the relevant higher-order structures that encode the causal relationships within a case are missing? Not every example is fully explained, when provided by an informant who assumes background knowledge, or when the example comes from observations. Another common problem in concept learning is when cases are dissimilar at the surface level or share confounding similarities. When a specific form of conclusion is available in a concept learning task, the Conclusion-Verified Analogical Schema Induction (CASI) method introduced here can be useful to overcome both these obstacles. CASI modifies how SAGE evaluates whether a mapping is a good candidate for generalization by relying on a consistency check that scrutinizes the mapping, rather than using the mapping's similarity score as a simple threshold. Though SAGE is a model of human analogical generalization and concept learning, CASI's modified algorithm is not presented as such a model.

The structure-mapping models used in this work are described in Research Background I begin the discussion here by introducing CASI, then describe two experiments evaluating it, and close with discussion and future work.

B. CONCLUSION-VERIFIED ANALOGICAL SCHEMA INDUCTION

SAGE's basic function relies on SME's computation of similarity to discern the proper basis around which to build generalizations: SME generates the mapping between the new case and the case retrieved by MAC/FAC, and if the mapping score is sufficiently high, the mapping is used to assimilate the new case into the existing generalization, or to form a new generalization if MAC/FAC retrieved an outlier. Analogy control predicates allow a certain amount of fine-tuned control over the mapping generated, for example by indicating that specific expressions in the base and target cases must map to each other, or that some expression in the case(s) should be mapped. There are circumstances, however, where SME may be led astray by distractor facts shared across cases. Furthermore, SME's similarity score is enhanced by shared higher-order structures within cases, i.e., statements that take other statements as arguments, such as causal relationships and constraining relations. When those higher-order structures are missing from the case representations, SME may assign cases a similarity score that is so low as to not be useful as a metric for determining whether cases should be assimilated. However, just because cases do not share the higher-order structures that are a significant contributor to SME's similarity score does not mean that no good schemas can be formed from such cases, only that SME may be unable to detect where those good schemas might come. Under those circumstances, if SAGE is to learn useful information it will require alternative criteria for when a mapping should be used for generalization.

Imagine one wants to learn about a principle and has a set of cases illustrating it. Crucially, the principle is unknown: the facts relevant to the principle are assumed to be present in the case, but the higher order structure tying some of those facts to the conclusions are missing. Each case has an outcome, and it is understood that the outcome is derived from the operation of the unknown principle over other facts of the case. The goal is to isolate the facts consistently associated with the outcomes, even if the specific principles that explain why those facts and outcomes are associated are unknown. By identifying those patterns of facts across a series of cases, the facts associated with particular conclusions—the facts that illustrate the unknown principle—can highlighted. The goal is to build generalizations that will allow the system to understand future cases relying upon that principle, cases whose conclusions—unlike those of the training cases—are unknown. This is the sense in which CASI does induction.[31]

The training cases, though potentially missing the explanation linking the facts to the outcomes, all have the information of what outcome was associated with which set of facts. I call these *solved cases*, with the particular piece of information that may be missing in future cases labeled as the *conclusion*, that can be differentiated from all other case facts. (The conclusion must have a form that can be prespecified to the system, so it can recognize conclusions in the cases where they appear.) The goal is to learn from solved cases to reliably infer that conclusion when appropriate.

Ordinarily one would trust SAGE to learn to extract the relationships between the problem facts and the conclusions. But when the higher-order relationships between the problem facts and

[31] CASI- and SAGE-style schema induction, wherein the mapping itself becomes the schema, differs from other models of schema induction—e.g., the one used by the LISA symbolic-connectionist model of analogical reasoning and schema induction, (Hummel & Holyoak, 2003)—where a schema is created by essentially elements of cases that participate in a mapping into a new case.

conclusions are not included in training cases, and especially when the cases contain distractingly similar but irrelevant facts, SAGE does not always produce reasonable answers. While there is always the risk of setting SAGE's assimilation threshold too high (leading to too few generalizations) or too low (leading to generalizations that do not capture useful information), I have found that in at least one dataset of cases there is no sweet spot where relying on SME's similarity score to control assimilation is enough to get useful generalizations. Instead, a mapping to be used for generalization must be examined to determine whether it in fact captures the information that is being used by SAGE to learn, before that mapping can be used for generalization.

To illustrate this consider our domain, Tort Law. Under the law of Trespass, a person is liable for trespassing on another's private property if the first person was on the property without permission or excuse. But imagine we do not know that principle. Instead, we are trying to build a schema of facts that illustrate the trespass rule, from cases that contain a set of facts and the conclusion of whether those facts encode a trespass. My dataset contains 29 cases where one party trespassed on the property of another; the cases in Table 5.1 are taken from my dataset and are typical of the cases in it.

Unfortunately, neither of these cases contain higher-order structures, so they cannot share any such structures. The required natural language syntax connecting the conclusion facts to the rest of the facts in the case are too complex to be expressible using the current capabilities of the

Table 5.2: Two cases from the dataset (case text simplified from original)

	Case 1: *Trout v Bank of Belleville*	Case 2: *Conklin v. Newman*
Case Text	The defendant is a bank. The bank has a parking lot. The parking lot's exit was blocked with a chain. The chain was hard to see at night. The plaintiff was riding his motorcycle at night. The plaintiff drove into the bank's parking lot. The plaintiff collided with the chain. The plaintiff died.	The plaintiff's father owned land. A fence divided the land into two pieces. The plaintiff lived on one piece of the land. The defendant lived on the other piece of the land. The fence was on both pieces of the land. Then the plaintiff's father gave her her piece of the land. Then the defendant destroyed the fence. Then the plaintiff's father died.
Predicate Logic Representations (extracted by CNLU)	(deathOf die68193 plaintiff67673) (defendants Trout_v_Bank_of_Belleville defendant66716) (degreeOfDifficulty see67334 (HighAmountFn LevelOfDifficulty)) (doneBy block67045 chain67121) (doneBy have66867 defendant66716) (driverActor ride67699 plaintiff67673) (eventOccursAt ride67699 night67454) (instrument-Generic block67045 chain67121) (instrument-Generic collide68054 chain67121) (isa block67045 ObstructionEvent) (isa chain67121 Chain) (isa collide68054 Collision) (isa defendant66716 BankOrganization) (isa defendant66716 Defendant) (isa die68193 Dying) (isa drive67891 TransportInvolvingADriver) (isa exit66971 Exit) (isa have66867 Possession) (isa motorcycle67743 Motorcycle) (isa night67454 Night) (isa parking-lot66896 ParkingLot) (isa plaintiff67673 Plaintiff) (isa ride67699 TransportInvolvingADriver) (isa see67334 VisualPerception) (isa Trout_v_Bank_of_Belleville CourtCase) (objectActedOn block67045 exit66971) (objectActedOn collide68054 chain67121) (objectActedOn have66867 parking-lot66896) (objectFoundInLocation plaintiff67673 parking-lot66896) (occursDuring see67334 night67454) (performedBy drive67891 plaintiff67673) (plaintiffs Trout_v_Bank_of_Belleville plaintiff67673) (possessiveRelation defendant66716 parking-lot66896) (possessiveRelation parking-lot66896 exit66971) (possessiveRelation plaintiff67673 motorcycle67743) (primaryObjectMoving collide68054 plaintiff67673) (to-Generic drive67891 parking-lot66896) (topicOfIndividual see67334 chain67121) (transporter ride67699 motorcycle67743)	(cardinality group-of-piece14864 2) (comesFrom-Generic piece15433 land14512) (deathOf die16855 father14430) (defendants Conklin_v_Newman defendant15339) (doneBy destroy17908 defendant15339) (doneBy divide14674 fence14618) (doneBy own14480 father14430) (elementOf piece14864 group-of-piece14864) (elementOf piece15433 group-of-piece14864) (eventOccursAt destroy17908 piece14864) (eventOccursAt destroy17908 piece15433) (father plaintiff14417 father14430) (fromAgent give18445 father14430) (isa Conklin_v_Newman CourtCase) (isa defendant15339 Defendant) (isa destroy17908 DestructionEvent) (isa destroy17908 (CausingFn DamageOutcome)) (isa die16855 Dying) (isa divide14674 Parcellation) (isa father14430 HumanFather) (isa fence14618 Fence) (isa give18445 GivingSomething) (isa group-of-piece14864 SetOfTypeFn SeparatePhysicalPartOfObject)) (isa land14512 RealEstate) (isa live15010 ResidingSomewhere) (isa live15347 ResidingSomewhere) (isa own14480 OwningSomething) (isa piece14864 SeparatePhysicalPartOfObject) (isa piece15433 SeparatePhysicalPartOfObject) (isa plaintiff14417 Plaintiff) (objectActedOn destroy17908 fence14618) (objectActedOn divide14674 land14512) (objectActedOn own14480 land14512) (objectGiven give18445 piece14864) (on-UnderspecifiedSurface fence14618 piece14864) (on-UnderspecifiedSurface fence14618 piece15433) (outputsCreated divide14674 group-of-piece14864) (owns plaintiff14417 piece14864) (parts land14512 piece14864) (parts land14512 piece15433) (plaintiffs Conklin_v_Newman plaintiff14417) (possessiveRelation plaintiff14417 father14430) (possessiveRelation plaintiff14417 piece14864) (residence-Role live15010 piece14864) (residence-Role live15347 piece15433) (residents-Role live15010 plaintiff14417) (residents-Role live15347 defendant15339) (startsAfterStartingOf destroy17908 live15010) (startsAfterStartingOf die16855 destroy17908) (startsAfterStartingOf give18445 live15347) (to-Generic divide14674 group-of-piece14864) (toAgent give18445 plaintiff14417)
Conc.	(trespassOnPropertyByAction plaintiff67673 parking-lot66896 drive67891)	(trespassOnPropertyByAction defendant15339 land14512 destroy17908)

language system that was used to translate the cases. Because of that lack of shared higher-order structure, SME generally scores the similarity of these cases quite low. In fact, if the assimilation threshold is set even to 0.1 (a very low score), only 3 of the 29 cases in the dataset will be assimilated together. And if it is set lower still (say, 0.01), generalizations are generated, but of poor quality and do not capture the information that defines a trespass and that is shared across many cases. There is no “sweet spot” between these scores. Why does this happen?

Examining these two cases, it should be clear to the reader that, from the perspective of trespass, the proper analogs here for the trespassers are the plaintiff in Case 1 and the defendant in Case 2, the proper analog for the property being trespassed upon are the parking lot and the land, and the trespassees should be the defendant in Case 1 and the plaintiff in Case 2. But SME gets distracted by the facts that both plaintiffs own something (in Case 1, a motorcycle, and in Case 2, a piece of land), and by the irrelevant shared statements about death. In particular, it gets distracted by multiple statements involving a relation (like objectActedOn and possessiveRelation) and aligns things it should not align (Figure 5.1). Because the trespass statements are not aligned, if this mapping is used to create a generalization, the system will have learned nothing about trespass.

Perhaps all that is needed is simply to require that the trespass conclusions (the statements describing the conclusion that a trespass occurred) correspond with each other? Unfortunately, that

Figure 5.1: A Mapping between Cases1 and 3 that is Useful for Generalization

Base Item (*Trout*)	**Target Item (*Falejczyk*)**
(Dying die68193)	(Dying die16855)
(objectActedOn have66867 parking-lot66896)	(objectActedOn destroy17908 fence14618)
(doneBy have66867 defendant66716)	(doneBy destroy17908 defendant15339)
(doneBy block67045 chain67121)	(doneBy own14480 father14430)
(objectActedOn block67045 exit66971)	(objectActedOn own14480 land14512)
(possessiveRelation plaintiff67673 motorcycle66743)	(possessiveRelation plaintiff14417 piece14864)

does not solve the problem in these cases: SME will put those two facts—and therefore the entities in the specific conclusion facts—in correspondence, without properly aligning the *other* statements involving those same entities that explain why those entities play the role they do in the conclusion statement (Figure 5.2). And having correctly aligned the plaintiff in Case 1 with the defendant in Case 2, but incorrectly aligned the Death statements with each other, it cannot even align the role relations attached to the Death statement, since doing so would violate SME's 1:1 mapping constraint (by aligning the plaintiff in Case 1 – who is already aligned with the defendant in Case 2 – with the father in Case 2). Thus this is a mapping that puts the conclusions into correspondence, but does not provide much useful information about trespass.

In fact, while these two cases were chosen for illustration purposes because SAGE consistently assimilates them together when using the lower assimilation threshold, they actually should *not* generalize together, because SME will not generate a useful mapping from them for learning about trespass. *Trout* should assimilate based on a mapping with *Falejczyk*, a third case concerning someone who paves over and parks on his neighbor's lawn (Figure 5.3), while my experiments suggest *Conklin* is an outlier that should not assimilate with other cases in the dataset (at least, not without rerepresenting the facts of the case). Unfortunately, the only way to prevent SAGE from generating a generalization from the mismatched cases is to set the assimilation

Figure 5.2: Mapping between Cases 1 and 2 with Required Conclusion Correspondence

Base Item (*Trout*)	**Target Item (*Conklin*)**
(Dying die68193)	(Dying die16855)
(trespassOnPropertyByAction plaintiff67673 parking-lot66896 drive67891)	(trespassOnPropertyByAction defendant15339 land14512 destroy17908)
(doneBy block67045 chain67121)	(doneBy divide14674 fence14618)
(objectActedOn have66867 parking-lot66896)	(objectActedOn own14480 land14512)
(doneBy have66867 defendant66716)	(doneBy own14480 father14430)

threshold sufficiently high that no other useful generalizations are made either. Thus, when similarity scores are sufficiently low because of a lack of higher-order structure shared across cases and because of shared distractor entities and relations, SME's similarity score becomes an ineffective standard by which to determine whether two cases should be assimilated. That holds true even when the relevant information (the case conclusions) are required to be mapped. In such a situation, will SAGE simply be unusable for datasets that have these qualities?

Fortunately, this is not the case. One can use a variant of SAGE to construct useful generalizations by ensuring that a task-specific conclusion predicate plays a very specific role in the mapping (described below). SME is perfectly capable of generating the mapping that both maps the conclusion and those relevant statements, but it might not get a chance to, either because a higher-scored mapping is used for generalization first, or (as I found in my experiment) the cases that *should* be usefully generalized together cannot be because they have already been assimilated with other cases in ways that occlude the useful information about those cases. In other words, in tasks where there is a known type of conclusion to be drawn when learning a concept, the algorithm

Figure 5.3: A Mapping between Cases1 and 3 that is Useful for Generalization

Base Item (*Trout*)	**Target Item (*Falejczyk*)**
(possessiveRelation defendant66716 parking-lot66896)	(possessiveRelation plaintiff3950 property3961)
(doneBy have66867 defendant66716)	(doneBy own3895 plaintiff3950)
(objectActedOn have66867 parking-lot66896)	(objectActedOn own3895 property3961)
(TransportInvolvingADriver drive67891)	(TransportInvolvingADriver park4128)
(to-Generic drive67891 parking-lot66896)	(to-Generic park4128 property3961)
(objectFoundInLocation plaintiff67673 parking-lot66896)	(objectFoundInLocation defendant3885 property3961)
(trespassOnPropertyByAction plaintiff67673 parking-lot66896 drive67891)	(trespassOnPropertyByAction defendant3885 property3961 park4128)
(objectActedOn block67045 exit66971)	(objectActedOn own3894 property3902)
(objectActedOn collide68054 chain67121)	(objectActedOn pave4688 part4736)

can use that constraint to search for the most productive prior generalizations or outliers for assimilating a new example. That is what CASI does.

Conclusion-verified Analogical Schema Induction (CASI) is essentially a consistency check that replaces SAGE's reliance on an assimilation threshold when learning in tasks where there is a known type of conclusion to be drawn. The algorithm is presented in Figure 5.4. CASI works by withholding the conclusion from the probe case (step 1) when performing an analogical retrieval for generalization (step 3). It then checks whether that conclusion is amongst the candidate inferences from the retrieved mapping (steps 4-7). That is, it checks whether the mapping from the retrieved (solved) case to the probe's case facts *without* the conclusion would allow SME to generate the probe's withheld conclusion. The held-out conclusion thus verifies the mapping. If the mapping generates the held-out conclusion, the conclusion is reintegrated into the probe case

Figure 5.4: the CASI Algorithm.

Given case *c*, gpool *g*, conclusion predicate *cp*:
CASI(*c*, *g*):
 1. Probe *pc* = nonConclusionFacts(*c*)
 2. Conclusions *SC* = {conclusion(*c*)} [each of form cp(X)]
 3. Reminding *r* = reminding(*pc*, *g*)
  4. **If** *r*:
   5. **For** mapping *m* in *r*:
    6. **For** *sc* ∈ *SC:*
     7. **If** *sc* ∈ candidate-inferences(*m*):
       8. RetrievedCase *retr* = baseOfMapping(*m*)
       9. doSageGeneralizeWithMapping(*c*, *retr*, *m*)
     Else: [do nothing; go to next conclusion]
    Else: [do nothing; go to next mapping]
   If no mapping:
    10. CASI(*c*, caseLibMinus(*g*, *retr*)
    11. [perform another reminding sans retrieved cases]
  Else: [no retrieved cases, add *c* ungeneralized]
    Else: [do nothing; go to next mapping]
   If no mapping:
    10. CASI(*c*, caseLibMinus(*g*, *retr*)
    11. [perform another reminding sans retrieved cases]
  Else: [no retrieved cases, add *c* ungeneralized]

and the mapping, and SAGE uses the mapping to assimilate the probe with the retrieved case (steps 8-9). (If there is more than one conclusion statement, CASI will use the first mapping that generates one of them.) Otherwise, it examines another mapping for the same retrieval, followed by additional retrievals, (steps 10-11) until it finds one that works or runs out of candidates.

The key idea is to ensure that the mapping used to generalize example cases together captures the connections between the case facts and the case conclusions, connections that might be implicit and only revealed across multiple cases governed by the same principle. This is using analogy for inductive reasoning: SME constructs a mapping that explains what is common across cases, and CASI verifies that the mapping in fact can explain that which the system is meant to learn. Thus CASI checks whether the facts SME has identified as shared can be used to project the conclusions to the cases in question. The reason this works is not because CASI necessarily forces SAGE to construct and use a good mapping, but rather that CASI will lead SAGE to *reject* assimilating cases with useless mappings, leaving those cases available to be assimilated with once the right future case (and mapping) comes along.

There is to my knowledge no evidence that humans engage in this kind of reasoning when learning concepts, though it is not implausible that they might. Humans can inspect the analogies they form between cases and the inferences they draw from those analogies, and they might well do so when they use those same analogies to learn about concepts. But CASI was developed to achieve more accurate performance on a machine learning task, not to implement a model of human cognition or make predictions about human performance. As such, my evaluations pit CASI against baseline SAGE and against a large-language model, rather than evaluating whether CASI can replicate results from experiments on humans.

C. EXPERIMENTAL VALIDATION

i. Experiment 1: Examining Generalizations

CASI was evaluated on Trespass, Assault, and Battery cases from the Illinois Intentional Tort Qualitative Dataset (Chapter 4:. This dataset of historical Illinois cases includes the original statement of case facts, syntactically simplified statements of those facts, machine translations of the simplified facts into predicate logic using the Companions Natural Language Understanding system (Tomai & Forbus, 2009), predicate logic representations of case conclusions, and information about the cases such as the decision year, the court, the legal claim at stake, and whether the claim was successful or not. A *positive* case is one where the legal claim is found to have occurred, i.e., one of the parties trespassed; a *negative* case is one where the claim failed. My experimental validation only used CASI to create generalizations from positive cases, because in the legal domain only positive cases should be expected to have relevant information in common. That is, positive cases are the ones that encode the events in which the claim at issue happened, and often negative cases only have in common the fact that they are not positive cases (unless there is a particularly common set of facts under which people repeatedly bring failing legal claims). While negative cases can be useful for delineating category boundaries (McLure et al., 2015), using negative cases in this way was not a part of my initial experiments for CASI. In all, CASI was tested on 29 positive cases in Trespass, 12 positive cases in Assault, and 30 positive cases in Battery.

CASI was first tested in comparison to the traditional method of using SAGE, with a match constraint to ensure that a legal case conclusion participated in the mapping. That is, the control condition involved feeding cases into SAGE as usual and letting SAGE pick the best mapping with which to generalize, provided the mapping included an expression correspondence that mapped the conclusion fact of the probe.[32] The experimental condition used the CASI algorithm as described above in Figure 5.3. Cases were assimilated within gpools[33] specific to the case doctrine (i.e., there was one gpool for Trespass cases, another for Assault cases, etc.). Cases were given to both algorithms in the same order, a hand-generated ordering designed to group like cases together. This manual ordering was designed to help SAGE and CASI find the best generalizations they could, to make those generalizations as useful as possible for automated legal reasoning (described in the subsequent two chapters, which also investigate using a random order of cases).

One of the motivations for developing CASI is that legal cases that feature the same claim may be extremely dissimilar, or similar in ways that can be distracting to SAGE. I therefore tested both algorithms using assimilation thresholds of 0.01 (a low standard of similarity reflecting the dissimilarity between cases)[34] and of 0 (relying on CASI's conclusion check in the experimental condition and the mapping constraint in the control condition rather than SME's similarity score).[35] CASI's consistency check is meant to replace reliance on SME's similarity score, and I indeed found that CASI using a threshold of 0 performed better than when using any threshold at all. Similarly, when using a threshold of 0, SAGE will assimilate all cases into a single generalization. A single generalization produces a joint probability table that can be used to produce Bayes nets

---

[32] I do not require that the solution statements in the probe and the retrieved case map to each other because some cases have more than one solution statement. SME's 1:1 mapping constraint prevents any statement in a case from corresponding to more than one statement in the other. When mapping a case with two solutions to a case with one, requiring solutions to correspond would involve arbitrarily deciding which solutions should correspond (or trying to break the 1:1 mapping constraint, thus producing no mapping). Requiring only that the probe's solution facts be mapped allows SME the flexibility of mapping as many such facts as it can, and if not all of them can be mapped, picking the best one(s) for the mapping.

[33] A gpool is the case library used for generalization, and which contains generalizations and ungeneralized outliers. See Chapter 2.D.

[34] Pilot investigation found that SME assigned almost no pairs of cases a normalized similarity score of even 0.1, leading me to use a very low similarity threshold in order to produce any generalizations at all.

[35] An assimilation threshold of zero has been used to produce joint probability tables for other ML algorithms (Halstead & Forbus, 2005).

or probabilistic rules (Halstead & Forbus, 2005), but does not generally produce a case that is useful for analogical reasoning. I therefore only report results of CASI using a threshold of 0 and SAGE using a threshold of 0.01. A probability cutoff of 0.6 was used, meaning that when a case is constructed for analogical reasoning from a generalization, facts below 0.6 probability are excluded from the case. I selected a probability cutoff of 0.6 as a sweet spot, high enough to include only facts present in a majority of cases in a generalization (to isolate the facts core to a legal claim), but low enough to include facts present in two of three cases in three-case generalizations.

The two approaches were evaluated by examining the generalizations generated by each approach in each doctrine. I looked at the facts in each generalization that were above the probability cutoff (i.e., the facts SME would put into a case constructed from that generalization) to see whether they contained the facts relevant to the legal doctrine at issue. To avoid the possibility of user bias, I pre-generated the facts necessary to find the legal claim. I verified whether those facts were in a given generalization, rather than make a judgment call as to whether the generalization was somehow "sufficient." The facts relevant to each doctrine are presented in Table 5.2. Note that the representation system uses neo-Davidsonian event representation, so the sentence "Pat walked into the house," for example, would be represented with three statements:

Table 5.2: Facts Sought In Each Generalization.

Doctrine	Facts Necessary to Infer Doctrine	# Necessary Facts
Trespass	(1) Trespassee owns property; (2) Trespasser is on property OR (a) Trespasser owns an object AND (b) the object is on the property; (3) The trespassing event is done by the trespasser OR by the trespasser's object; (4) The trespassing event brings the trespasser or the trespasser's object onto the property.	4-5
Assault	(1) A threat occurred; (2) The threat was performed by the assaulter; (3) The threat was against the assaultee.	3
Battery	(1) A touch occurred; (2) The touch was offensive OR the touch was harmful; (3) The batterer performed the touch; (4) The victim was the person touched.	4

Table 5.3: Comparing CASI and SAGE Generalizations.

Method	Doctrine	# Genl'ns Produced	# Ungenl'd Cases	% Genl'd Cases	# Correct Genl'ns	# Partial Genl'ns	# False Genl'ns	Proportion Correct Genl'ns	Proportion Partial+ Genl'ns
CASI	Trespass	9	5	83	4	2	3	0.444	0.667
	Assault	3	3	75	2	1	0	0.667	1
	Battery	6	6	80	2	2	2	0.333	0.667
SAGE	Trespass	5	14	52	0	1	4	0	0.2
	Assault	3	2	83	2	0	1	0.667	0.667
	Battery	6	12	60	2	1	3	0.333	0.5

one defining a walking event, one identifying Pat as the doer of that event, and one defining the event as being into the location of the house (plus statements declaring the house to be a house, and Pat to be a human).

Method	Doctrine	# Genl'ns Produced	# Ungenl'd Cases	% Genl'd Cases	# Correct Genl'ns	# Partial Genl'ns	# False Genl'ns	Proportion Correct Genl'ns	Proportion Partial+ Genl'ns
CASI	Trespass	9	5	83	4	2	3	0.444	0.667
	Assault	3	3	75	2	1	0	0.667	1
	Battery	6	6	80	2	2	2	0.333	0.667
SAGE	Trespass	5	14	52	0	1	4	0	0.2
	Assault	3	2	83	2	0	1	0.667	0.667
	Battery	6	12	60	2	1	3	0.333	0.5

Generalizations were coded into three categories: a "correct" generalization contained all the necessary facts; a "partial" generalization contained all but one necessary fact; and a "false" generalization was missing more than one necessary fact. The data are reported in Table 5.3.

All of the cases were selected to be relevant to illustrating the particular doctrine they represent, and for being similar – at least in my opinion – to other cases. Thus each case could be expected to usefully participate in a generalization of the legal principle in question. Therefore the percentage of generalized cases can be seen as a sort of measure of recall: the algorithm's effectiveness at finding and making use of useful cases. On the other hand, the ultimate question is how good the generated generalizations are. The proportion of correct generalizations (and similarly of correct + partially correct generalizations) can therefore be seen as a measure of precision: the proportion of learned concepts that were useful.

Method	Doctrine	# Genl'ns Produced	# Ungenl'd Cases	% Genl'd Cases	# Correct Genl'ns	# Partial Genl'ns	# False Genl'ns	Proportion Correct Genl'ns	Proportion Partial+ Genl'ns
CASI	Trespass	9	5	83	4	2	3	0.444	0.667
	Assault	3	3	75	2	1	0	0.667	1
	Battery	6	6	80	2	2	2	0.333	0.667
SAGE	Trespass	5	14	52	0	1	4	0	0.2
	Assault	3	2	83	2	0	1	0.667	0.667
	Battery	6	12	60	2	1	3	0.333	0.5

CASI outperforms SAGE on trespass cases. Indeed, SAGE was essentially unable to learn useful generalizations from trespass cases. Moreover, the proportion of generalizations that are

Method	Doctrine	# Genl'ns Produced	# Ungenl'd Cases	% Genl'd Cases	# Correct Genl'ns	# Partial Genl'ns	# False Genl'ns	Proportion Correct Genl'ns	Proportion Partial+ Genl'ns

correct *or* partially correct using CASI substantially exceeds those using SAGE. And closer examination revealed that the missing high-probability statement in many of the partial generalizations using CASI was the role relation indicating the perpetrator of the action. An examination of those generalizations revealed that these role relations often differed across cases (e.g., doneBy, performedBy, bodilyDoer, etc.). If these statements were re-represented to reflect their shared nature (i.e., they connect an action to its actor), I hypothesize that many partial CASI generalizations would become correct generalizations. On the other hand, SAGE generalizations in this domain often contained high-probability facts that had nothing to do with the outcomes of cases, for example, that the parties in both cases were walking dogs or driving a car at the time the tortious behavior occurred. Experiment 1 suggests that in this domain, CASI will synthesize more cases together and in a better way, resulting in more generalizations that will be potentially usable for reasoning by other systems, such as the ones described in this thesis. Whether this is the case is the subject of Experiment 2.

ii. Experiment 2: Using Generalizations

Constructing generalizations from cases is not done for its own sake, but so that learned concepts can be used for some task. I examined the performance of a legal reasoning system that reasons about legal cases using analogical generalizations, comparing its performance when using SAGE versus CASI generalizations. These legal reasoning techniques are described and evaluated in greater detail in the following chapters. I tested two of the experimental legal reasoning systems developed for this thesis—Analogical Reasoning with Positive Generalizations and Reasoning with Rules Learned from Generalizations—and compared them to a Large-Language Model (LLM) baseline. Both experimental systems involve holding out a case from the dataset, constructing generalizations from the remaining positive legal cases (using either SAGE or CASI), then using the resulting generalizations to reason about the held-out case.

In Analogical Reasoning with Positive Generalizations (ARPG), the generalizations are applied to the held-out case directly by analogy. As mentioned previously, in the legal domain cases illustrating the same legal principle can be highly dissimilar to each other except for the facts that define that legal principle. For reasons discussed in Analogical Generalization, Reasoning, and Rule Learning for Legal Reasoning about Common Law Torts., I assume that when learning from such dissimilar cases, properly formed generalizations ought only to encode those facts that define the claims at issue, with facts incident to individual cases having low probabilities and therefore not participating in the case constructed from the schema for analogical reasoning.[36] ARPG relies on this idea by examining the candidate inferences generated by a mapping from the generalization to the held-out case: if there is only one candidate inference, in the form of a conclusion statement, then all other facts in the generalization—and thus core to the legal claim—participate in the mapping and therefore have a corresponding fact in the held-out case. If the held-out case contains all facts core to the legal claim (other than the legal conclusion), then the held-out case contains an instance of that legal claim being met. On the other hand, if there are candidate inferences in addition to one for the legal conclusion, then there are facts present in the generalization that are not present in the held-out case. These missing facts correspond to what legal facts would be required for the held-out case to be a positive example of the legal claim in

[36] The assumption that generalization will strip away cases' idiosyncratic facts is specific to the legal domain, where concepts underlying legal doctrines are fairly abstract and can be grounded in a wide variety of different specifics. I do not assume in general that a schema-building process will strip away *all* facts incident to cases illustrating some concept. Cases contain all sorts of facts that may be correlated with a case's outcomes, and if enough cases in a generalization share those correlates, then the generalization will as well. A larger discussion of why I believe the assumption is reasonable in the legal domain is in Analogical Generalization, Reasoning, and Rule Learning for Legal Reasoning about Common Law Torts., where the approaches are described in detail.

question. ARPG thus functions by mapping a legal schema onto a new case and examining whether any of the schema facts are unmapped, and one way to do so is by seeing whether those facts are projected as candidate inference. So ARPG concludes whether a legal case is positive or negative simply by counting the extra candidate inferences: if the only candidate inference is the legal conclusion, then ARPG concludes that the case is a positive instance; if there are extra inferences besides the conclusion, it concludes that the case is a negative instance.[37] (ARPG only reasons with generalizations because ungeneralized cases will contain extra facts and therefore generate many candidate inferences.)

I tested ARPG using Precision@6, meaning the system checked that it generated the correct answer in the first six mappings it tried.[38] I did this to separate the system's ability to generate the proper answer from its ability to do so using the first mapping from its first retrieval, that is, to separate the system's ability to reason with a case from its ability to find the right case with which to reason on the first try. (Given the low SME similarity scores when comparing cases in this dataset, the retrieval task poses its own problems and is its own potential research area.) I also ran several conditions varying the number of additional candidate inferences ARPG would tolerate before concluding that a case was negative. The claim that legal generalizations should encode only facts relevant to the legal claim at issue is a theoretical postulation, and the results of

[37] Again, extra candidate inferences might simply correspond to additional correlate facts, and not be evidence that some concept is inapplicable. By hypothesis, the legal domain is an exception: because the facts common across legal cases in some domain operate at a fairly high level of abstraction, they are less situationally-specific and will therefore will share fewer correlates. As a corollary, any high-probability facts that survive in a sufficiently large generalization will also generally be present in the cases being reasoned about, and therefore will not be proposed as candidate inferences. Thus if a sufficiently large legal generalization includes facts other than the con-sequent or its antecedents, those correlative facts should generally also be in the new case being reasoned about, and so will not impede ARPG's performance. See the following chapter for a deeper discussion of this issue.

[38] The number 6 was chosen somewhat haphazardly: in setting up the GPT-J baseline, the code that I adapted automatically had it generate three completions, and I ran it twice. Having six completions for each case using GPT-J, I decided to examine all six and to test ARPG and PAPR using the same number.

experiment 1 demonstrated that these generalizations do not yet perfectly encode the legal claims at issue. The generalizations often contained extra facts incidental to the legal claim that I theorize would fall away with more cases assimilated, but that did not in my experiments. Therefore I ran ARPG with a tolerance for 0, 1, and 2 extra candidate inferences, to compare its performance on SAGE and CASI generalizations. Allowing extra candidate inferences acknowledges that there may be facts incidental to a claim that participate in a generalization.

I also compared SAGE and CASI generalizations using the technique described in Chapter 6: called Reasoning with Rules Learned from Generalizations (RRLG). RRLG constructs Horn clauses from generalizations and uses those rules to reason about a case using backchaining. RRLG constructs its Horn clauses by first replacing the generalized entities in the generalization with logical variables. It then extracts the conclusion statement and installs it as the Horn clause's consequent (conclusion predicates are known, so conclusion statements are identifiable), and installs the other facts of the generalization as antecedents of the Horn clause. RRLG filters out Horn clauses whose antecedents will not bind all consequent variables. RRLG's performance is evaluated by checking that it can correctly derive legal conclusions in positive cases, and correctly *fail* to derive a legal conclusion in negative cases.

It may be noted that I am measuring the performance of ARPG using Precision@6 while measuring RRLG's performance using Precision@1 (because RRLG only generates a single answer). For this experiment, I compared the performance of RRLG and ARPG to *themselves* when reasoning with legal schemas created using SAGE versus when using CASI, not to each other. The goal was to evaluate the best possible performance of each system using CASI versus SAGE generalizations, to best reveal the difference in performance that can be attributed to the generalizations themselves, and ARPG performed better using Precision@6. (This is discussed in the subsequent chapter, where these reasoning techniques' performance are also compared to each other.) I also compared their performance to that of legalBERT, a LLM BERT model specialized on legal cases. LegalBERT was not used to develop legal schemas (or trained on those schemas); it was tested instead by turning each case into a multiple-choice question by varying the original case conclusion (reversing the parties' roles, reversing the legal outcome, and reversing both the parties and the outcome). Its performance was assessed by prompting it with the simplified text description of a case's facts and checking which answer it selected. These results are presented in Table 5.4.

Table 5.4: Results (Experiment 2): Cases correctly solved by each approach, in absolute terms and as percentages, and Precision. ARPG algorithms tested using Accuracy@6.

Technique	Method	Overall Cases	Assault Cases	Battery Cases	Trespass Cases	Positive Cases	Negative Cases	Pos Prec.	Neg Prec.
ARPG 0CIs	CASI	35 (35%)	6 (35%)	10 (25%)	19 (44%)	7 (10%)	28 (97%)	.778	.308
ARPG 0CIs	SAGE	31 (31%)	6 (35%)	12 (30%)	13 (30%)	4 (6%)	27 (93%)	.571	.290
ARPG 1CIs	CASI	47 (47%)	9 (53%)	16 (40%)	22 (51%)	19 (27%)	28 (97%)	.905	.354
ARPG 1CIs	SAGE	33 (33%)	7 (41%)	12 (30%)	14 (33%)	10 (14%)	23 (79%)	.588	.277
ARPG 2CIs	CASI	58 (58%)	9 (53%)	22 (55%)	27 (63%)	30 (42%)	28 (97%)	.938	.412
ARPG 2CIs	SAGE	36 (36%)	9 (53%)	22 (55%)	13 (30%)	10 (14%)	26 (90%)	.714	.302
RRLG	CASI	47 (47%)	8 (47%)	21 (53%)	18 (42%)	23 (33%)	24 (83%)	.793	.338
RRLG	SAGE	48 (48%)	10 (59%)	24 (60%)	14 (33%)	19 (27%)	29(100%)	.95	.363
legalBERT	-	33 (33%)	9 (53%)	14 (36%)	10 (23%)	23 (33%)	10 (35%)	.535	.175

Results were compared using proportion tests (all significance results reported at $p<0.05$). ARPG with 0 additional CIs performed statistically the same with both CASI and SAGE forming the generalizations used. However, with 1 or 2 additional CIs tolerated, ARPG using CASI generalizations performed significantly better than when using SAGE generalizations. Notably, these techniques using CASI generalizations not only performed significantly better overall, but also specifically on Trespass cases. This is consistent with the results from Experiment 1 showing that CASI made the greatest improvement in schema learning for Trespass cases.

RRLG did not perform significantly differently when using CASI generalizations than when using SAGE generalizations. However, there was a non-significant trend of RRLG performing better on Trespass cases and on Positive cases when using CASI than when using SAGE, and of performing better on Assault and Battery cases when using SAGE than when using CASI. RRLG performed significantly better on negative cases when using SAGE than when using CASI (indeed, when using SAGE, RRLG got all the negative cases correct).

The trend in RRLG's results, though not significant, are consistent with the claim that CASI leads to better legal generalizations from this dataset. Consider RRLG's evaluation: if a case is positive, then to solve the case, RRLG must successfully fire one of the rules it learned from a generalization of other cases in the same legal doctrine. But when a case is negative, successfully solving the case means that all of RRLG's learned rules *failed* to fire. Perversely, this means that the worse RRLG's rules are at encoding some legal doctrine, the *better* it will perform in close negative cases, even as it performs worse on positive cases. Thus, the fact that RRLG using CASI's performance trends better on positive cases than when using SAGE, despite (and consistent with) its worse performance on negative cases, is an encouraging signal that CASI's generalizations are actually capturing relevant information about the legal principles governing these cases.

The observation that the quality of learned legal principles is better revealed by performance on positive cases than on negative ones holds true for ARPG as well. And in ARPG with 1 or 2 extra CIs, performance is not only significantly improved on Trespass cases when using CASI relative to SAGE, but on positive cases as well.

Even as performance was better on negative cases, precision was higher for positive cases: the systems were more likely to be correct when identifying a positive case than a negative case. For ARPG, precision was higher on both negative and positive cases when using CASI generalizations than SAGE generalizations. For RRLG, precision was higher (and performance trended better) when using SAGE generalizations.

Comparing my techniques to the baseline, RRLG when using both CASI and SAGE significantly outperformed legalBERT. ARPG when using SAGE was never able to significantly outperform legalBERT. When using CASI, ARPG with 0 extra CIs allowed did not significantly outperform legalBERT, but allowing 1 extra CI allowed ARPG using CASI to significantly outperform legalBERT (as did allowing 2 extra CIs).

D. DISCUSSION, LIMITATIONS, & FUTURE WORK

These results demonstrate that Conclusion-verified Analogical Schema Induction can be an effective tool for learning generalizations of facts underlying concepts. The experiments demonstrate that while CASI cannot guarantee that perfect generalizations are formed from my dataset, the generalizations formed capture useful conceptual information about the domain. More importantly, where CASI had the greatest effect improving generalizations in Experiment 1 (i.e., on Trespass cases), that improvement was reflected in improved performance when reasoning with those generalizations in experiment 2.

CASI is similar to Inductive Logic Programming (Muggleton & De Raedt, 1994) in that it is a form of inductive inference that operates over symbolic representations. Both involve examples that start with irrelevant information which is stripped away. ILP always produces rules, whereas analogical generalization produces probabilistic schemas and maintains outliers, both of which are applied via analogy, rather than unification. ILP typically operates offline, in batch mode, whereas analogical generalization is incremental. ILP uses both positive and negative examples in crafting its rules, to maximize coverage of positive examples and minimize coverage of negative examples, whereas CASI currently only uses positive examples. While SAGE has been ex-tended to incorporate automatically-derived near-misses, and thereby benefit from negative examples (McLure, Friedman, & Forbus, 2015), extending CASI to use near-misses is an avenue for future work.

CASI can only be used if the learning system knows ahead of time the specific form of the conclusion facts within its training data, because CASI requires removing the conclusion from a case during the first step generalization. CASI has been shown to be useful when learning from case sets that give SME little to grasp onto in guiding generalization, for example, when causal information connecting case facts to outcomes is unknown, or in domains where cases may be dissimilar to each other or share irrelevant distractor features. Legal reasoning is such a domain because cases illustrating legal principles may be dissimilar in every way except for the specific facts directly relevant to the purportedly illegal conduct. It remains to be seen whether CASI leads to improved learning and performance in areas where SAGE is already an effective learning system.

The twin observations that CASI's performance relative to SAGE is more improved for Trespass cases than for Assault and Battery, and that CASI performs better with a threshold of 0 than 0.01, supports the conclusion that CASI is most useful for use in datasets where SME will consistently assign very low similarity scores to case comparisons. CASI's improved performance with a threshold of 0 suggests that even the vanishingly low assimilation threshold is an impediment to generalizing useful cases in this domain. Furthermore, because of the nature of the dataset and how it was generated, Assault and Battery cases contained more information specific to legal claims than did Trespass cases. That is, Assault cases consistently represent legally relevant information with the predicates threateningAgent and threatenedAgent, while legally relevant information in Battery consistently involves an entity that is a TouchingEvent. In contrast, the facts describing legally-relevant information in Trespass cases use role relations used to represent many different kinds of events in the cases. These common and repeated role relations make cases distractingly similar to each other even as similarity scores remained low. For datasets such as the one presented here, SME's similarity score is ineffective at signaling to SAGE when two cases should assimilate.

CASI has a clear limitation relative to SAGE, which is in its efficiency. SAGE efficiently uses MAC/FAC and SME to find mappings, and if the mappings score high enough, produces a generalization. Using CASI requires scrutinizing each mapping's candidate inferences, and potentially going back to the well repeatedly until a mapping either produces the required conclusion or all cases in the library have been exhausted (although it would be trivial to set a maximum number of retrievals CASI could perform).

This efficiency limitation is related to the fact that SAGE is a model of how humans naturally learn concepts through comparison of similar cases, one that reproduces human psychological results (Forbus, Ferguson, Lovett, & Gentner, 2017), while CASI represents a substantially more deliberative, conclusion-oriented problem-solving approach and is not a model of everyday human reasoning. I was inspired by my domain: lawyers not only use analogies, but legal reasoning and argument is more deliberative and rigorous than the everyday learning that humans naturally engage in. That said, while it may be the case that humans at times engage in such cognitively intensive conclusion-deriving analogical reasoning for concept learning, I currently have no evidence that they do. Determining whether CASI in fact tracks with deliberative human cognition is one potential area of future work; another is to test CASI on more cases and more domains. It remains to be seen whether CASI will only lead to improved performance in complex domains featuring highly dissimilar cases such as legal reasoning, or whether it is a useful technique for concept learning in general.

Now that I have described the dataset and the technique developed to learn the concepts encoded therein, I will turn to my experimental legal reasoning techniques and my evaluations thereof.

CHAPTER 6: ANALOGICAL GENERALIZATION, REASONING, AND RULE LEARNING FOR LEGAL REASONING ABOUT COMMON LAW TORTS.

This chapter concerns the three legal precedential reasoning algorithms I developed. Two are analogical reasoning techniques used to bring the generalizations learned by CASI (described in the prior chapter) to bear on held-out cases in the dataset. The third is a method to turn the learned analogical generalizations of legal rules into actual logical rules, i.e., Horn clauses. The chapter will describe the algorithms and an evaluation thereof.

Before describing the systems, I want to make a note about where this research fits into the legal scholarship concerning common law reasoning. As described in Research Background the legal scholarly academy does not agree about the role analogies play in legal reasoning. I do not claim that my AI research sheds light on how human lawyers engage in precedential reasoning. However, I believe that when human lawyers are searching through a library of precedents they do not forget each precedent they have examined before moving on to the next one. Because humans naturally recognize analogies and engage in a process of analogical generalization when examining similar cases in sequence, I believe that to the extent that legal precedents encode rules, analogical learning is the process by which those rules are revealed. Furthermore, if this theory is correct, there is an intermediate state during which an analogical schema has been created through the comparison of those precedent cases, before it is turned into a rule. Those schematized cases—those legal principles—can also be applied by analogy to the case at bar.

Learning rules directly from a range of cases also addresses a persistent problem in common law case law: that the rule declared in any given case may not be the rule that actually governs the broader area of law. That is, a judge at Time 1 might declare, "The rule is {A, B, C, D} → X. Factors A, B, C, and D are present, and therefore outcome X is required by law." Then at Time 2, a judge confronts a case with only factors A, B, and C, but is quite confident that outcome X is appropriate; the judge might then declare something like "The prior court may have *seemed* to say that the rule is {A, B, C, D} → X, but it's clear to any reader that the real rule is {A, B, C} → X, and factor D just made it even worse." Learning the rules by examining what actually happened across a series of cases, rather than only taking a single statement of the rule as dispositive, could lead to more accurate understanding of what actually governs some legal domain.

I thus present three algorithms for legal reasoning, each of which begins by using analogical generalization to learn schemas of legal precedents. The first algorithm resolves cases directly by analogy, both to the learned schemas and to any ungeneralized cases that did not participate in any schema. The second algorithm creates analogies between learned legal schemas and a case at bar, then reasons about the constructed analogy—as opposed to reasoning only *with* the analogy—in order to reason about the case at bar. The third algorithm takes the learned schemas and converts them into rules, reasoning about the case at bar using deductive logic. I thus provide demonstrations not only of how cases can be resolved directly by analogy or using rules, I also present a mechanism by which legal principles or rules can be learned from the precedent cases that govern some doctrine.

A. THREE ALGORITHMS FOR PRECEDENTIAL LEGAL REASONING

All three algorithms presented here use CASI to build generalizations from cases in a given domain and reason about held-out cases in light of what the system has learned through generalization. These algorithms contribute to both AI & Law research and to a debate in the legal academy by describing a framework for the role and nature of analogy in legal learning and reasoning. Taken together, they describe a cycle by which legal rules are discerned and applied from a body of precedent cases.

The first algorithm is Purely Analogical Precedential Reasoning (PAPR), which involves using SME to reason by direct analogy about held-out cases in light of the learned generalizations and ungeneralized cases. The second algorithm, Analogical Reasoning with Positive Generalizations (ARPG), also involves using SME to reason about held-out cases, but only in light of positive generalizations, not negative generalizations or ungeneralized examples. ARPG resolves cases using meta-reasoning about the constructed analogies, rather than directly by analogy. The third algorithm, Reasoning with Rules Learned from Generalizations (RRLG), transforms the positive generalizations used in ARPG into Horn clauses, and applies those rules to held-out cases using back-chaining. Although ARPG and RRLG were introduced in the prior chapter, they are reintroduced here with more detail on their motivation and function.

i. Purely Analogical Precedential Reasoning

The first algorithm, Purely Analogical Precedential Reasoning (PAPR), is presented in Figure 6.1. PAPR first gathers all cases in the dataset in the same doctrine as the case at bar and creates generalizations from them using CASI.[39] (The legal doctrine at issue is always known;

Figure 6.1: Purely Analogical Precedential Reasoning (PAPR) Pseudocode

Given case *c*, case set *cs*, conclusion predicate *cp*:
PAPR(*c*, *cs*):
- 1. *pos* = positiveCases(*cs*); *neg* = negativeCases (*cs*)
- 2. Gpools *pg* = CASIgeneralize(*pos*); *ng* = CASIgeneralize(*neg*)
- 3. Probe *pc* = nonConclusionFacts(*c*)
- 4. Gpool Case Library *cl* = union(*pg*, *ng*)
- 5. Reminding *r* = reminding(*pc*, *cl*)
- 6. **If** *r*:
    - 7. **For** mapping *m* in *r*:
        - 8. Inferences *CIs* = candidateInferences(*m*)
        - 9. *groundCIs* = {*i* in *CIs* if not containsSkolem(*i*)}
        - 10. **If** *groundCIs*:
            - 11. **For** *inf* ∈ *groundCIs*:
                - 12. **If** predicate(*inf*) = *cp*:
                    - 13. **Return** *inf* as conclusion to *c*
        - 14. **Else**: [if here, the top retrieved mapping failed to solve the case, try others]
- 15. **Else**: [if here, no mapping in *r* provided a solution; strike rejected cases & try again]
    - 16. Retrieved cases *retr* = retrievedCases(*r*)
    - 17. Updated Case Library *cl* = *cl* – *retr*
    - 18. [go to step 5].

Note: when using Precision@N testing, replace step 13 with the following steps tracking N and validating the conclusion:
- 13(a). *depth* += 1 [*depth* is initialized to 0]
- 13(b). **If** *depth* > *N*: **Return** fail. [N reached]
- 13(c). Conclusions *SC* = {conclusion(*c*)}
- 13(d). **If** *inf* ∈ *SC*: **Return** *inf*

---

[39] All my methods were evaluated using hold-one-out training and testing. I did this because of the relatively small number of cases, some of which only had one or two analogs: I wanted to be sure that those analogs would always be present in the training set. With a larger dataset using cross-fold validation would be a better approach.

even without knowing if a party is liable, one can know what they are accused of). One gpool is used for positive cases (where a party was eventually found liable) and another for negative ones.

After generalizations are generated, PAPR performs a reminding over the union of positive and negative gpools using the case at bar as a probe, and gathers the generated mappings from the retrieved cases to the case at bar. The candidate inferences from the top mapping that represent conclusion statements (i.e., hypothesized solutions) are inspected, and inferences that contain skolem variables—candidate inferences that involve hypothesizing new entities rather than explaining the case using those already present—are rejected. The algorithm then returns as the case outcome the skolem-free conclusion candidate inference from its top-scored mapping. To separate the algorithm's capacity to get the correct answer from its ability to get it right the first time, I also ran a condition where the experimental system was able to check its answer against the held-out truth and examine other mappings or perform additional retrievals. In this latter condition the system examines up to six mappings that generate skolem-free conclusion inferences, giving a measure of Precision@6.[40]

ii. Analogical Reasoning with Positive Generalizations

The second algorithm, Analogical Reasoning with Positive Generalizations (ARPG), reflects the fact that negative cases are united largely by not being positive cases rather than by anything they have in common with each other. That is, two positive trespass cases may both feature a person on someone else's private property without permission, but the only thing that might unite two negative trespasses is the *absence* of such facts. Thus, while negative cases might help define the boundaries of a legal generalization by showing what a legal claim does *not* involve

[40] Again, the number six was chosen because my initial run of GPT-J happened to generate six completions, so I used this number for consistency.

(using, for example, a system like the one described in McLure et al. 2015), it is the positive cases that demonstrate what a legal claim is. ARPG exclusively reasons with reference to positive generalizations, to measure the extent to which the learned schemas distill the facts relevant to a legal claim. Like PAPR, ARPG looks for a mapping with a grounded case conclusion candidate inference. Once it has one, ARPG examines whether there are also any *other* candidate inferences, and the presence or absence of such other candidate inferences forms a critical part of its functionality. ARPG's current implementation thus depends on an important and unproven hypothesis, which requires some explanation.

A HYPOTHESIZED PROPERTY OF LEGAL CASE GENERALIZATIONS

Consider what legal generalizations contain: nothing but the facts common to the cases that participate in the generalization. I hypothesize that legal cases are so varied that a sufficiently broad variety of positive legal cases should *only* have in common the facts relevant to the legal claim and facts that will be true of almost all such claims, with the facts idiosyncratic to particular cases but irrelevant to the claim stripped away through the process of building schemas.[41] In turn, this hypothesis implies that, in general, an analogy between a learned legal schema and a case that contains all the elements of a legal claim will generate only one candidate inference (for the case conclusion), because all the other facts in the schema will have a corresponding counterpart in the case being reasoned about.

Before I explain why I believe this, I will note that the hypothesis being wrong would not be fatal to ARPG's approach to reasoning about cases. ARPG's method relies on having a case

[41] This is even more true when learning from judicial opinions, where judges have decided to include or omit facts from the record for a variety of reasons, doing much of the work of determining which information is legally relevant. Facts that are always true but irrelevant to legal doctrine are not likely to be described in the opinions, and are therefore less likely to survive into generalizations.

that contains only the facts core to a legal claim with which to reason by analogy. My hypothesis is that the analogical generalization process is sufficient to generate such a case, but perhaps I am mistaken.[42] If so, it does not mean that ARPG would not be a viable approach to legal reasoning, but only that some process other than or in addition to the analogical generalization process would be necessary to generate those cases for reasoning.

Why would I believe that the generalization process would be sufficient? After all, cases contain all sorts of facts that may be correlated with a case's outcomes, and if enough of the cases in a generalization share those correlates, then the generalization will as well. If so, when using that generalization to reason about a new case, one should assume that before long one will be reasoning about a case that happens not to contain the correlate fact, and that the correlate fact will be proposed as a candidate inference. The presence of extra candidate inferences should therefore not, in general, be taken as evidence of the inapplicability of the primary concept learned by the generalization, because they may well simply respond to those correlate facts, and not just the antecedents to the concept.

I believe the legal domain is an exception to the general rule that concepts learned from generalizations will be accompanied by all sorts of candidate inferences irrelevant to the case being reasoned about. Because the facts common across legal cases in some domain operate at a fairly high level of abstraction, they are less situationally-specific and will therefore will have fewer correlative facts in common *across* a variety of cases. For example, an assault is defined as a credible threat of immediate harm: although many cases may involve weapons, others will involve

[42] My experimental results, at least, are insufficient to affirm or reject the hypothesis: my evaluations do not yield analogical generalizations that contain only the facts core to a legal claim, but this might be attributed to relatively small size of the dataset or the fact that the cases are represented without the deeper shared structures that best support analogical learning.

none; many cases will involve the threat of physical damage, but others will involve the threat of, for example, inappropriate touching (which causes a different kind of harm than a violent strike). As the facts specific to individual legal cases are stripped away through the generalization process (for example, that a particular weapon was used, or the nature of the harm threatened), the correlates of those specific facts will also be stripped away with them. What is left are the facts core to the legal claims, and correlates of those core facts. And certainly such correlates may exist, either because they are associated with the specific facts core to the legal claim—for example, that the threatening person was looking at the person being threatened—or because they are so common that they are true in general—for example, that the assaultor was clothed. If these correlates are sufficiently common, they may be true even in most legal cases and may therefore remain in the schema.

But if correlates remain in the generalization, will they not be proposed as candidate inferences? No, because if those correlates are so common as to be present in the generalization of legal concepts, they should also generally be present in the cases being reasoned about. Even if looking at someone is not necessary to credibly threaten to imminently harm them, in most assault cases the assaultor *will* look at the assaultee. Even though a nude person is capable of assaulting one another (and certainly must have at various points in time), the vast majority of assault cases will involve a clothed attacker. Because these facts will be present in the cases being reasoned about, they will therefore not be proposed as candidate inferences, and their absence will not be remarked upon.

I recognize that this hypothesis is just that, and it may well be wrong. And admittedly, when presented for the first time with the rare case of a nude person committing an assault, ARPG might conclude that an assault did not occur based on the additional candidate inference about the assaultor being nude. This is a substantial failure mode, and is one of the reasons I do not suggest that ARPG should be directly entrusted with legal decision-making. Trimming these excess facts from a generalization when faced with a counterexample is an area of future work, along with determining how to know when a generalization has assimilated enough cases that it is reasonable to make the assumption that low-probability facts have fallen away, and how to rerepresent case facts at the right level of abstraction for generalization. But the fact that one can expect a particular failure mode in edge cases does not mean that the technique is useless for reasoning about cases in the general case.

Thus, while this may not be true in other domains, in theory schemas generated from a sufficiently varied set of solved legal cases should include only abstracted legally relevant facts, the legal conclusion those facts lead to, and correlates that are likely to also be true in new cases that encode the legal principle. Put differently, the generalizations should thus be largely self-contained legal principles, associating a legal conclusion with an abstract set of facts that lead to and are associated with that conclusion. And analogizing an unsolved legal case—with all the elements of a legal claim, and by assumption the correlates of those elements—to that generalization should therefore generate only one candidate inference: the case conclusion.

But again, even if my hypothesis is wrong, it does not mean that ARPG's approach to reasoning about cases is useless, but only that ARPG cannot rely solely on the analogical generalization process to generate the schemas with which it reasons. Some other process would need to be developed to remove extraneous facts from generalizations. Such a process would not only be useful for the legal domain, but for concept learning in general.

With this assumption clearly defined and justified, I turn to the specific ARPG algorithm, which is presented in Figure 6.2.

THE ANALOGICAL REASONING WITH POSITIVE GENERALIZATIONS ALGORITHM

ARPG starts by creating generalization just as PAPR does, but this time it only generalizes the positive cases in a given doctrine (again excluding the case at bar, if it is a positive case).[43] ARPG then discards ungeneralized examples, leaving only the positive generalizations themselves.[44] It then begins reasoning by proceeding as PAPR does, using the case at bar's problem as a probe to retrieve over the generalizations, and looking for candidate inferences that take the form of a case solution (and that do not include skolem variables).

ARPG takes the mappings with grounded case conclusions in its inferences and examines how many *other* candidate inferences the mapping has. If it has other candidate inferences, ARPG concludes that the case is a negative case, and that the other candidate inferences represent the missing elements of the claim. If the conclusion statement is the only candidate inference, ARPG concludes it is a positive case.

ARPG has an additional partial truth check in the form of tolerating extra candidate inferences, i.e., how many extra candidate inferences besides the case conclusion can be tolerated before concluding that the case is a negative example. Tolerating extra candidate inferences allows the system to be able to reason with generalizations that may not have stripped away all legally

---

[43] Negative cases are useful sources of information about what must *not* be true in order to conclude that a case is positive: permission, legal standing, etc. Such cases can be used to sharpen category boundaries in analogical generalization (McLure, Friedman, & Forbus, 2015), but that is left to future work in this thesis.

[44] At the point at which ungeneralized examples are discarded, what ARPG has learned will only be used for reasoning, not for subsequent learning. Those ungeneralized examples might still be useful for learning if new cases come in that might generalize with them. To reflect this possibility, my implementation of ARPG copies the learned generalizations over into a new reasoning context, leaving the original gpool untouched and able to continue learning if new cases were to come in.

irrelevant facts. As with PAPR, in addition to evaluating the ARPG method on its first returned answer, I evaluated it using Precision@6. In that condition, if ARPG is either mistaken about what

Figure 6.2: Analogical Reasoning with Positive Generalizations (ARPG) Pseudocode

Given case *c*, case set *cs*, conclusion predicate *cp*, number of extra candidate inferences tolerated *error*:

ARPG(*c*, *cs*):

1. *pos* = positiveCases(*cs*)
2. gpool of all cases *gAll* = CASIgeneralize(*pos*)
3. reasoning gpool *g* = generalizationsInGpool(*gAll*)
4. Probe *pc* = nonConclusionFacts(*c*)
5. Reminding *r* = reminding(*pc*, *g*)
6. **If** *r*:
    7. **For** mapping *m* in *r*:
        8. Inferences *CIs* = candidateInferences(*m*)
        9. *groundCIs* = {*i* in *CIs* if not containsSkolem(*i*)}
        10. **If** *groundCIs*:
            11. **For** *inf* ∈ *groundCIs*:
                12. **If** predicate(*inf*) = *cp*:
                    13. *otherCIs* = *CIs* - *inf*
                    14. **If** count(*otherCIs*) > *error*: [If there are more non-conclusion CIs than the tolerated *error*, this is a negative case]
                        15. **Return** (not *inf*) as conclusion to *c*
                    16. **Else**: **Return** *inf* as conclusion to *c*
        17. **Else**: [if here, the top retrieved mapping failed to solve the case, try others]
18. **Else**: [if here, no mapping in *r* provided a solution; strike rejected cases and try again]
    19. Retrieved cases *retr* = retrievedCases(*r*)
    20. Updated Case Library *g* = *g* - *retr*
    21. [go to step 5].

Note: when using Precision@N testing, replace steps 13–16 with the following steps tracking N and validating the conclusion:

13. *depth* += 1 [*depth* is initialized to 0]
14. **If** *depth* > *N*: **Return** fail. [N reached]
15. *otherCIs* = *CIs* – *inf*
16. **If** count(*otherCIs*) > *error*:
    17. **If** isNegativeCase(*c*):
        18. **Return** (not *inf*) as conclusion to *c*
    19. **Else**: [Wrong answer, keep trying]
20. **Else**: [The conclusion CI is the only CI]
21. Conclusions *SC* = {conclusion(*c*)}
22. **If** *inf* ∈ *SC*: Return *inf*

kind of case (positive or negative) the case at bar is, or if it correctly concludes that it is a positive case but generates the wrong conclusion, it will move on to other mappings or perform a re-retrieval, examining up to six mappings with grounded case conclusions.

Where PAPR's performance can be properly understood to measure the extent to which analogy can be used to solve legal cases (by analogy to other cases and generalizations), ARPG is more a measure of the extent to which the system producing the schemas for reasoning has learned accurate and useful legal principles that can be used to solve cases (here, the CASI modification to the SAGE analogical generalization system). ARPG simply will not work if it does not have sufficiently clean generalizations of legal principles. Thus PAPR is a measure of the usability of analogy as a technique to reason about legal cases, whereas ARPG is a measure of the effectiveness of analogical generalization in learning legal concepts (but still assesses analogy's utility in applying those learned concepts).

In this dataset, each solution takes the form of a ternary predicate, expressing who did what to whom. For example, `(assaultsPartyByDoing Fred Rick punch123)` means that `Fred` assaulted `Rick` when he performed the action `punch123`; similarly, `(not (trespassOnPropertyByAction Carl lawn456 walk789))` means `Carl` did not trespass on the property `lawn456` by taking action `walk789`. The structure of these answers means PAPR and ARPG can also use a *partial truth check* to measure of the *extent* to which a mapping produces the right answer. The partial truth check requires the first argument in the ternary predicate to be correct (i.e., the accused must be correctly identified), but then is satisfied with only one of the two remaining arguments. That is, it would rate `(assaultsPartyByDoing Fred Rick kick246)` and (not `(trespassOnPropertyByAction Carl house357 walk789))` as partially true statements of the above conclusions.

iii. Reasoning with Rules Learned from Generalizations

The final algorithm, Reasoning with Rules Learned from Generalizations (RRLG), is presented in Figure 6.3. RRLG closes the loop on how legal *rules* (not just legal *principles*) are extracted from precedent cases and applied to a case at bar, by converting ARPG's positive generalizations into rules and using them to reason about a new case. RRLG's algorithm by which generalizations are converted into rules depends on the fact that each generalization contains a high-probability fact representing the generalized legal principle around which the generalization is constructed (for example, `(assaultsPartyByDoing <generalizedEntity1> <generalizedEntity2> <generalizedAction3>)`).

RRLG begins with the same generalizations generated for ARPG, and proceeds as follows. First, RRLG discards all facts below a certain probability threshold, preserving the same facts that

Figure 6.3: Reasoning with Rules Learned from Generalizations (RRLG) Pseudocode

Given case *c*, case set *cs*, conclusion predicate *cp*, probcutoff *p*:
RRLG(*c*, *cs*):
1. *ruleset* = {}
2. *pos* = positiveCases(*cs*)
3. gpool of all cases *rawG* = CASIgeneralize(*pos*)
4. reasoning gpool *g* = *rawG* - outliers(*rawG*)
5. **For** *genl* ∈ *g*:
    6. *highProbFacts* = {*f* for *f* in genlFacts(*genl*) where P(*f*)>*p*}
    7. *genEntVars* = generalizedEntities(*genl*)
    8. *logicVars* = {(*i*, logicVar(*i*))} for *i* in *genEntVars*
    9. *rawConc* = conclusion(*genl*) [of form *cp*(X)]
    10. *rawAntes* = {*highProbFacts* – *rawConc*}
    11. *ruleConc* = replaceVars(*rawConc*, *logicVars*)
    12. *antes* = {replaceVars(*a*, *logicVars*)} for *a* in *rawAntes*
    13. *Horn* = makeHornClause(*ruleConc*, *antes*).
    14. **If** ∀ variables(*ruleConc*) ∈ variables(*antes*):
        15. *ruleset* += *Horn*
16. Case facts *cf* = nonConclusionFacts(*c*)
17. Conclusions *Concs* = {conclusion(*c*)}
18. **For** *conc* ∈ *Concs:* [query the case conclusions given case facts and learned rules]
    19. *ans* = query(*conc*, *cf* + *ruleset*)
    20. **If** *ans***: Return** *ans*
21. **Else**: **Return** "*c* is a negative case" [no conclusions derived]

would be put into a generalization's case for reasoning by analogy (my experiments used a probability cutoff value of 0.6). Next, RRLG converts all generalized entities into logical variables. Although it never came up with these cases, RRLG does not replace ungeneralized entities with variables, since any ungeneralized entity participating in a high-probability fact is involved in several cases, and therefore presumably important to the principle governing them. Next RRLG produces a Horn clause, installing the variablized legal conclusion[45] as the consequent of the rule and all other facts as antecedents. Finally, RRLG examines each rule to ensure that every variable in the consequent has a counterpart in the antecedents, to ensure no rule will leave a consequent variable unbound.[46]

To apply RRLG to a case at bar, the algorithm queries for the conclusion of the case at bar in a query context that shares the facts of the case and the rules learned by RRLG. For positive cases, it labels the case as correct if it is able to derive the correct legal conclusion; for negative cases, it labels the case as correct if it cannot derive the legal conclusion.

These three algorithms define a spectrum regarding the role of analogy in legal reasoning. All three algorithms use analogical generalization to find principles common to legal cases, but how those principles are used varies by technique. In PAPR the algorithm seeks to resolve cases purely by analogy: the generalizations might improve analogical reasoning, but PAPR treats generalized and ungeneralized—and positive and negative—cases as equally informative. ARPG still reasons by analogy, but restricts the learned principles, and reasons *about* the analogies it has drawn between those principles and the case at bar, not just *using* those analogies. Finally, in

[45] Identifiable by its predicate.

[46] The fact that this was necessary suggests an efficiency improvement to PAPR and ARPG, to filter generalizations whose high-probability facts do not ground all the entities in the case conclusion. This might reduce the number of retrievals that those methods rejected for having conclusion inferences containing skolem variables.

RRLG the case at bar is not reasoned about by analogy at all, but using rules extracted from the learned legal principles.

B. EXPERIMENTAL VALIDATION

I tested my methods against each other and several off-the-shelf machine learning baselines, on Assault, Trespass, and Battery cases from the Illinois Intentional Tort Qualitative Dataset. The experimental dataset includes 17 assault cases (12 positive, 5 negative), 40 battery cases (30 positive, 10 negative), and 43 trespass cases (29 positive, 14 negative). The greater number of positive than negative cases reflects the fact that judges are more likely to publish case opinions in cases where a party is found liable (i.e., cases that illustrate what the law is and how it applies) than those where they were not (i.e., cases that illustrate what the law is not).

I ran RRLG, PAPR, ARPG, and several baselines on these cases, varying PAPR and ARPG's parameters. First, I ran both PAPR and ARPG using both the strict truth test (i.e., did the algorithm generate the entire legal conclusion as a candidate inference) and using the partial truth test (i.e., did the legal conclusion candidate inference include the correct alleged wrongdoer and either the correct victim or the correct tortious action). I also ran ARPG with a tolerance for either 0, 1, or 2 extra candidate inferences allowed when inferring a positive case, to determine whether the generalizations generated might be noisy. Finally, I tested PAPR and ARPG using both Precision@1 and Precision@6 testing. For Precision@1 the system would return the first answer it generated with a grounded conclusion inference, but for Precision@6, it would check its answer against the held-out ground truth: if it was wrong and the system had not yet checked six mappings, it would move on to other mappings and retrievals. (RRLG can only generate one answer, so Precision@6 testing is impossible.)

SAGE, like humans, is sensitive to the order in which it receives cases. I therefore tested both PAPR and ARPG on generalizations from a randomized case order (randomized for each new case at bar, i.e., not the same random order across cases), and on generalizations from a hand-selected case order that grouped like cases together for input into CASI. Because I found that a non-random order was more effective, I only ran RRLG using the hand-selected case order. Data on the generalizations generated and used for reasoning are provided in Table 6.1. Recall that each algorithm is run using hold-one-out testing, so the specific generalizations generated varied across runs, although not from method to method, since all relied on CASI. These data illustrate how few generalizations tended to be formed in negative cases, with sometimes none being formed at all.

For baselines, I used two BERT models and two GPT-based models. All models were retrieved from HuggingFace's model library and were tested on the simplified English descriptions of the cases that had been used to generate the predicate logic representations operated over in my experimental conditions.

Table 6.1: CASI Generalizations Generated for Legal Reasoning Experiment

Doctrine	Valence	Average # of Generalizations	Average # of Outliers	Max # Generalizations	Min # Generalizations
Trespass	Positive	8.6	5.3	9	7
	Negative	2	8.8	2	2
Assault	Positive	2.9	2.3	4	2
	Negative	0.9	2.1	1	0
Battery	Positive	5.7	6.3	6	4
	Negative	1	7.7	1	1
Total	Positive	6.5	5.3	9	2
	Negative	1.4	7.2	2	0

The BERT models were roBERTa and legalBERT, a version of roBERTa pretrained on legal text. To test the BERT models I turned the cases into multiple-choice tests and had the model select the best answer. Three false solutions to each case were generated by taking the original case solution and (1) reversing the parties' roles (i.e., "the plaintiff assaulted the defendant" became "the defendant assaulted the plaintiff"), (2) negating the original conclusion ("the plaintiff did not assault the defendant"), and (3) negating and reversing the conclusion ("the defendant did not assault the plaintiff"). Like RRLG, the BERT models always score the same choices the same way, so Precision@N testing is meaningless.

The GPT models were a public release of GPT-2 and GPT-J, a model based on GPT-3 with 1.3 Billion parameters. I fine-tuned and tested the GPT models using 5-fold cross-validation, each time training on 80% of the dataset's cases' problems and solutions, and testing on the remaining 20% of the case problems. I tested the GPT models by using the case problem as a prompt and having the model generate up to 100 tokens as a completion, six different times (i.e., I generated 6 continuations for each case, once again testing with Precision@6). I then read the case completions to determine whether they contained the correct solution to the case, both looking only at the first answer generated (Precision@1) and at all six (Precision@6). For GPT-J I defined an answer as partially true if it contained the correct solution regardless of whether it also contradicted itself, and strictly true if it did not contradict itself.[47] For example, if the correct solution was "the plaintiff assaulted the defendant," a GPT-J answer that included the phrases "the plaintiff assaulted the defendant. The plaintiff did not assault the defendant." would be partially true but not strictly true. All told I thus had 17 variations of the experimental conditions and 5

[47] I did not bother examining partial truth for GPT-2 because GPT-J so dramatically outperformed it.

baseline techniques; with both Precision@1 (first-answer) and Precision@6 testing, I examined a total of 32 conditions.

Techniques were compared using proportion tests. As an initial matter, the analysis demonstrated that legalBERT and GPT-J (strict truth test) significantly outperformed roBERTa and GPT-2 (both $p < 0.0001$). Furthermore, when directly comparing my own experimental methods (PAPR and ARPG) on randomized vs. non-randomized training sets, I found that using my hand-selected nonrandom order consistently trended better than using random order, but never significantly outperformed the random order (either overall or when looking at individual doctrines or case valences). To simplify reporting results (and to report only the higher-performing version of each method), I therefore only report results on the legalBERT and GPT-J baselines, and only on the nonrandom training orders of my own techniques (because I developed RRLG after having performed initial tests on PAPR and ARPG, I did not test RRLG on random-ordered cases). The reported techniques are identified in Table 6.2. Table 6.3 reports results from each technique, presenting performance by raw score and percentage accuracy, overall and broken down by doctrine and case valence, and using Precision@1 and Precision@6. Head-to-head comparisons of techniques are presented in Tables 6.4 and 6.5, which present results using Precision@1 and Precision@6, respectively.

i. Precision@1 v. Precision@6

Before comparing the methods' performance, I offer two general observations regarding the evaluation. As noted above, I tested methods capable of generating multiple answers using both Precision@1 (taking the top-scored answer – here, the first answer generated – as the system's output) and Precision@6 (having the system check its answer against the held-out truth and try again if it was wrong, up to six times, which is equivalent to checking if the answer is in the top six system outputs) to separate a method's ability to generate the right answer from its ability to generate the right answer on the first try. As Table 6.3 demonstrates, evaluating using Precision@6 leads to improvement over using Precision@1 across the board, on both my own methods and on the GPT-J baselines. In fact, the only method where the improvement from testing with Precision@6 was non-significant was ARPG with 0 additional CIs allowed.

But interestingly, it was not only these methods' performance relative to themselves that changed, but also their performance relative to each other. Most obviously, GPT-J significantly underperformed nearly all other methods using Precision@1, even as it *outperformed* many methods using Precision@6. But more subtle differences appear. For example, when both were

Table 6.2: Reasoning techniques (sans roBERTa, GPT-2, or PAPR & ARPG w/ random training)

Technique Name	Algorithm	Training Order	Truth Check (Strict/Partial)	# Extra CIs Allowed?	Precision@6 Testing?
RRLG	RRLG	Nonrandom	Strict Truth	-	No
PAPR-NR-ST	PAPR	Nonrandom	Strict Truth	-	Yes
PAPR-NR-PT	PAPR	Nonrandom	Partial Truth	-	Yes
ARPG-NR-ST-0CIs	ARPG	Nonrandom	Strict Truth	0	Yes
ARPG-NR-ST-1Cis	ARPG	Nonrandom	Strict Truth	1	Yes
ARPG-NR-ST-2CIs	ARPG	Nonrandom	Strict Truth	2	Yes
ARPG-NR-PT-0CIs	ARPG	Nonrandom	Partial Truth	0	Yes
ARPG-NR-PT-1Cis	ARPG	Nonrandom	Partial Truth	1	Yes
ARPG-NR-PT-2CIs	ARPG	Nonrandom	Partial Truth	2	Yes
legalBERT	l-BERT	Random	Strict Truth	-	No
GPT-J-ST	GPT-J	Random	Strict Truth	No contradictions	Yes
GPT-J-PT	GPT-J	Random	Partial Truth	Any	Yes

evaluated using Precision@1, ARPG with a strict truth test and 0 extra CIs significantly outperformed PAPR with a strict truth test; when evaluated using Precision@6, PAPR outperformed ARPG. These results suggest that, for the analogical reasoning methods, performance can be improved not only by improving the reasoning itself, but by finding ways to retrieve the right case the first time around.

Comparing RRLG's performance using Precision@1 to the performance of other systems using Precision@6 may seem like comparing apples to oranges, but is useful for determining how these systems' relative performance might change as obstacles to their performance—like retrieval—are addressed. That is, RRLG does not depend on retrieval at all, and the mechanics of the retrieval system used in PAPR and ARPG are not theoretical commitments to how those

Table 6.3: Performance of each technique overall, by legal doctrine, and on positive vs. negative examples, with raw scores and percentage accuracy.

	Overall		**Assault**		**Battery**		**Trespass**		**Positive**		**Negative**	
Precision / **Technique**	P@1 # (%)	P@6 # (%)	P@1 # (%)	P@6 # (%)	P@1 # (%)	P@6 # (%)	P@1 # (%)	P@6 # (%)	P@1 # (%)	P@6 # (%)	P@1 # (%)	P@6 # (%)
RRLG	47 (47%)	-	8 (47%)	-	21 (53%)	-	18 (42%)	-	23 (33%)	-	24 (83%)	-
PAPR-NR-ST	17 (17%)	46 (46%)	3 (18%)	11 (65%)	9 (23%)	20 (50%)	5 (12%)	15 (35%)	14 (20%)	38 (54%)	3 (10%)	8 (28%)
PAPR-NR-PT	28 (28%)	72 (72%)	6 (35%)	16 (94%)	11 (28%)	27 (68%)	11 (26%)	29 (67%)	20 (28%)	53 (75%)	8 (28%)	19 (66%)
ARPG-NR-ST-0CIs	28 (28%)	35 (35%)	5 (29%)	6 (35%)	9 (23%)	10 (25%)	14 (33%)	19 (44%)	2 (3%)	7 (10%)	26 (90%)	28 (97%)
ARPG-NR-ST-1Cis	29 (29%)	47 (47%)	6 (35%)	9 (53%)	9 (23%)	16 (40%)	14 (33%)	22 (51%)	4 (6%)	19 (27%)	25 (86%)	28 (97%)
ARPG-NR-ST-2CIs	32 (32%)	58 (58%)	6 (35%)	9 (53%)	13 (33%)	22 (55%)	13 (30%)	27 (63%)	11 (16%)	30 (42%)	21 (72%)	28 (97%)
ARPG-NR-PT-0CIs	32 (32%)	43 (43%)	5 (29%)	7 (41%)	10 (25%)	12 (30%)	17 (40%)	24 (56%)	6 (9%)	15 (21%)	26 (90%)	28 (97%)
ARPG-NR-PT-1Cis	36 (36%)	58 (58%)	6 (35%)	10 (59%)	12 (30%)	20 (50%)	18 (42%)	28 (65%)	11 (16%)	30 (42%)	25 (86%)	28 (97%)
ARPG-NR-PT-2CIs	41 (41%)	71 (71%)	6 (35%)	10 (59%)	15 (38%)	26 (65%)	20 (47%)	35 (81%)	20 (28%)	43 (61%)	21 (72%)	28 (97%)
legalBERT	33 (33%)	-	9 (53%)	-	14 (36%)	-	10 (23%)	-	23 (33%)	-	10 (35%)	-
GPT-J-ST	12 (12%)	52 (52%)	1 (6%)	6 (35%)	3 (8%)	20 (50%)	8 (19%)	26 (60%)	11 (16%)	44 (62%)	1 (3%)	8 (28%)
GPT-J-PT	12 (12%)	61 (61%)	1 (6%)	6 (35%)	3 (8%)	23 (58%)	8 (19%)	32 (74%)	11 (16%)	52 (73%)	1 (3%)	9 (31%)

systems function; comparing RRLG to the analogical reasoning approaches using Precision@6 thus gives a sense of the algorithms' relative performance if retrieval was a solved problem.

ii. Strict Truth vs. Partial Truth

As with Precision@6, the partial truth test and the tolerance for additional candidate inferences when using ARPG were designed to examine whether the legal reasoning methods partially captured the relevant information in the cases about which they were reasoning. I therefore expected these relaxed standards to lead to improved performance, and they did. PAPR performed significantly better with a partial truth test than a strict one regardless of Precision@1 or @6, both overall and when broken down by doctrine. When using Precision@1 ARPG's performance improved when *both* the partial truth test was used *and* additional CIs were tolerated, but while each relaxed standard used alone trended better, those trends were mostly nonsignificant. However, when evaluating using Precision@6 it was more of a mixed bag: ARPG sometimes

Table 6.4: Performance of reasoning techniques compared, Precision@1. Cell values indicate which technique significantly outperformed the other, if any.

Technique	GPT-J-PT P@1	GPT-J-ST P@1	legalBERT	ARPG PT 2CIs P@1	ARPG PT 1CIs P@1	ARPG PT 0CIs P@1	ARPG ST 2CIs P@1	ARPG ST 1CIs P@1	ARPG ST 0CIs P@1	PAPR PT P@1	PAPR ST P@1
RRLG	RRLG	RRLG	RRLG	---	---	RRLG	RRLG	RRLG	RRLG	RRLG	RRLG
PAPR-ST P@1	---	---	BERT	ARPG	ARPG	ARPG	ARPG	ARPG	ARPG	PT	
PAPR-PT P@1	PAPR	PAPR	---	ARPG	---	---	---	---	---		
ARPG-ST-0CIs P@1	ARPG	ARPG	---	PT 2CIs	---	---	---	---			
ARPG-ST-1CIs P@1	ARPG	ARPG	---	PT 2CIs	---	---	---				
ARPG-ST-2CIs P@1	ARPG	ARPG	---	---	---	---					
ARPG-PT-0CIs P@1	ARPG	ARPG	---	---	---						
ARPG-PT-1CIs P@1	ARPG	ARPG	---	---							
ARPG-PT-2CIs P@1	ARPG	ARPG	---								
legalBERT	BERT	BERT									
GPT-J-ST P@1	---										

improved with a partial truth test alone, sometimes only by increasing the number of CIs tolerated, and sometimes both. Notably ARPG with 2 extra CIs and strict truth test outperformed ARPG with 0 CIs and a partial truth test. This suggests that the tolerance of additional CIs had more to do with the improved performance than the partial truth test, suggesting that focusing on stripping away extra facts that are still present in those generalizations is an important area of future work if ARPG is going to be a useful legal reasoning technique.

I also assessed partial truth in GPT-J, but it made no difference when testing with Precision@1, and only a non-significant trend of improvement when testing with Precision@6.

iii. Comparing Methods

I begin by comparing the reasoning techniques to each other, then describe how they fare compared to the baselines. (All measures reported as significant are at $p < 0.05$, and most are at p

Table 6.5: Performance of reasoning techniques compared, Precision@6. Cell values indicate which technique significantly outperformed the other, if any.

Technique	GPT-J-PT P@6	GPT-J-ST P@6	legalBERT	ARPG PT 2CIs P@6	ARPG PT 1CIs P@6	ARPG PT 0CIs P@6	ARPG ST 2CIs P@6	ARPG ST 1CIs P@6	ARPG ST 0CIs P@6	PAPR PT P@6	PAPR ST P@6
RRLG	GPT	---	RRLG	ARPG	---	---	---	---	RRLG	PAPR	---
PAPR-ST P@6	GPT	---	PAPR	ARPG	ARPG	---	ARPG	---	PAPR	PT	
PAPR-PT P@6	PAPR	PAPR	PAPR	---	PAPR	PAPR	PAPR	PAPR	PAPR		
ARPG-ST-0CIs P@6	GPT	GPT	---	PT 2CIs	PT 1CIs	---	2CIs	1 CIs			
ARPG-ST-1CIs P@6	GPT	---	ARPG	PT 2CIs	---	---	---				
ARPG-ST-2CIs P@6	---	---	ARPG	PT	---	ST 2CIs					
ARPG-PT-0CIs P@6	GPT	---	---	2CIs	1CIs						
ARPG-PT-1Cis P@6	---	---	ARPG	2CIs							
ARPG-PT-2CIs P@6	---	ARPG	ARPG								
legalBERT	GPT	GPT									
GPT-J-ST P@6	---										

< 0.001; a detailed breakdown is provided in Appendix B.) These comparisons are in Tables 6.3 and 6.4, which report results using Precision@1 and Precision@6, respectively.

The strictest measure for learning and reasoning using analogy is ARPG using a strict truth test and allowing no extra candidate inferences in a mapping to conclude that a case is a positive example (ARPG-ST-0CIs). Indeed, this was the lowest-scoring of my experimental approaches, both when testing using Precision@1 and Precision@6. When testing using Precision@1 there was little improvement from allowing additional candidate inferences or using a partial truth test, but when using Precision@6, allowing one or two additional candidate inferences created a significant improvement over allowing zero extra CIs. The improvement in performance from allowing one extra CI to allowing 2 was not significant, at $p=0.059$, although nearly as many additional cases were solved with each additional CI permitted.

When using Precision@1 and a strict truth test, ARPG consistently outperformed PAPR. As to using a partial truth test, when using Precision@1 PAPR and ARPG performed substantially the same. However, when using a partial truth test and evaluating using Precision@6, PAPR outperformed nearly every other condition, including most ARPG conditions. From this I conclude that, without any help or second chances, ARPG is more able to identify and leverage legally relevant information than PAPR, but only under circumstances where both methods' performance leaves significant room for improvement. But the dramatic improvement of both approaches when using Precision@6 and partial truth tests demonstrates that these methods are capable of capturing legally relevant information, just not all of it and not on the first try.

There are two differences between PAPR and ARPG that potentially confound the source of the difference in their performance: that PAPR reasons with both negative cases and with outliers, whereas ARPG reasons with neither. PAPR heavily relied on these outliers: Of the 17 cases that PAPR correctly solved using a strict truth test and evaluated using Precision@1, 11 were solved by analogy to an ungeneralized outlier (9/14 positives, 2/3 negatives). Though not a majority, of the 72 cases that PAPR correctly solved using a partial truth test and evaluated using Precision@6, 30 were solved by analogy to an outlier (18/53 positives, 12/19 negatives). Furthermore, ARPG's superior performance was driven by its high performance on negative cases; PAPR outperformed it on positive cases. These results demonstrate that ungeneralized cases remain an important resource for reasoning directly by analogy. This may especially be true for negative cases, since ungeneralized exemplars were used in the majority of negative cases that PAPR solved.

RRLG significantly outperformed nearly every other condition when I evaluated our analogical methods and GPT-J using Precision@1, but the improved performance largely disappeared when using Precision@6, and RRLG was outperformed by methods using partial truth tests. When using Precision@6, RRLG performed significantly better than ARPG-ST-0CIs, and on par with PAPR-ST. RRLG consistently outperformed the analogy techniques on negative cases and was outperformed on positive cases. However, ARPG-PT-2CIs and PAPR-PT—techniques with relaxed standards for concluding cases were correct—both significantly outperformed RRLG. Because performance using Precision@1 is equivalent to real-world performance on unseen cases where checking answers is impossible and partial truth does not count, I would expect RRLG to be the highest-performing method if these algorithms as they currently exist were deployed in the real world.

As to the baselines: when the system could only generate one answer (i.e., using Precision@1 testing), RRLG beat all baselines, all ARPG conditions beat the GPT-J baselines, and

PAPR beat the GPT-J baselines when using a partial truth test. When using a strict truth test, PAPR did not perform significantly differently from GPT-J, and was outperformed by BERT. Using Precision@6, however, was again a mixed bag: GPT-J with a partial truth test outperformed most other conditions, although GPT-J with a strict truth test only beat the strictest ARPG condition; it was beaten by PAPR using a partial truth test and ARPG using a partial truth test and tolerating 2 additional CIs, and did not otherwise perform significantly differently from my methods. Additionally, all my methods except ARPG with 0 extra CIs tolerated (strict or partial truth test) significantly outperformed legalBERT. ARPG with 0 extra CIs outperformed legalBERT on Trespass cases and negative cases, while legalBERT did better on positive cases. However, loosening ARPG's standard of correctness led to significant improvement over legalBERT. GPT-J's substantially improved performance using Precision@6 led it to close the gap with RRLG when using a strict truth test, and to outperform RRLG when using a partial truth test.

The two baselines' relative performance changed when evaluating GPT-J using Precision@1 vs. Precision@6. When using Precision@1, legalBERT beat GPT-J using strict and partial truth test. When using Precision@6, both GPT-J methods beat legalBERT.

Finally, comparing techniques' performance to themselves, ARPG and RRLG consistently do better on negative than positive cases; PAPR does better on positive cases (although not when using a partial truth check), GPT-J does better on positive cases, and legalBERT does not significantly outperform itself on one type of case over another.

C. DISCUSSION

There are several interesting things to note from these results. Most obviously, the improvement in my analogical reasoning techniques and in GPT-J when testing using Precision@6 demonstrates that these algorithms are currently more capable of generating correct answers than they are of generating them the first time around. That is, those methods often eventually get the right answer, but often not on the first try. This improved performance cannot be accounted for as the system simply exhausting all its possible mappings: when I ran my techniques with no depth limit to see what they could solve with unlimited tries and feedback, PAPR would regularly examine over a dozen mappings, and at times derived the correct answer after doing so. These methods' strong performance using Precision@6 testing therefore shows that good mappings are being retrieved and reasoned about in the first few cases, just not always as *the* first case. These results suggest that one of the most fruitful directions for future work in this area will come from focusing on the retrieval system, either by affecting what cases get returned first, or by investigating how retrieved cases can be validated before being used to reason about some case at bar.

The improvement in both PAPR and ARPG when using looser standards of correctness suggest that these algorithms are indeed capturing information about what governs and how to solve legal claims, but that the learned generalizations involved are still noisy, and research into further refining those generalizations will be useful. In particular, it may well be the case that the assumption underlying this particular implementation of ARPG—that the analogical generalization process alone will strip away irrelevant facts in the legal domain—is false. (This evaluation is insufficient to conclude whether the hypothesis is true or false regardless of the outcome.) If the hypothesis is false, that does not meant that the *reasoning* component of ARPG's method is useful, but that the CASI schema-building process representing the *learning* component is not by itself sufficient to generate clean schemas, and the schemas will have to be modified by some other process to strip those additional irrelevant facts away. If ARPG would be improved by improving generalizations, RRLG (which uses the same generalizations) should also be improved.

At nearly 50% accuracy, RRLG is already quite high-performing, and is my highest-performing method when testing using Precision@1. In the absence of a separate process that might select a system's later or lower-ranked output over its first or highest-ranked one, using a system with Precision@1 is equivalent to using it when the true outcome of the case is unknown, so I expect RRLG to be the best-performing condition on truly new cases. One possible explanation for RRLG's improved performance when compared to analogy techniques using Precision@1 but not using Precision@6 is that RRLG, unlike analogy, actually *does* exhaustively search through all learned knowledge, by firing all the rules derived from its schemas and seeing if any of them work. The analogical reasoning techniques must instead use the schemas one at a time, and so can be stymied if they pick the wrong one with which to reason first.

Also, note that these research techniques involve *generating* an answer rather than selecting one: the system must say who did what to whom, not simply pick whether a party is liable. From this perspective, the fact that my techniques outperform legalBERT, which is selecting amongst answers, is promising. Though GPT-J performed commensurately with PAPR (strict truth) and ARPG with extra CIs permitted when using Precision@6, had I required GPT-J to generate only *true* statements or to label events as our methods did, then almost none of its responses would be correct (only 11 of the over 600 generated completions were entirely true, and only one of those was generated on the system's first try). Also, since most training cases are positive, with the defendant as the accused, it is generally a good guess when using this dataset that the defendant behaved tortiously towards the plaintiff. This may explain why GPT-J performed better on positive cases. It also might explain why my reasoning techniques generally outperformed GPT-J on negative cases; another explanation is that ARPG's and RRLG's performance on negative cases may be inflated, because they correctly solve negative cases by failing to derive positive conclusions, and they might perform worse on negative cases as they do better on positive ones.

I want to draw special attention to the techniques involving converting positive generalizations to rules and those involving reasoning with positive generalizations with a strict truth test. These are the techniques that most require the algorithm to understand the actual principles resolving precedent cases, and to apply those principles to new cases. My rule-learning technique represents a model of how human lawyers might reason about cases: by extracting rules from precedent cases and applying those rules to a case at bar. While I do not claim to provide evidence that this is in fact how human lawyers learn legal rules, I believe it provides a plausible account.

The ARPG algorithm is able to generate the correct answer in a third of cases even with the strictest standard of understanding applied. As the standard loosens, the performance goes up. And unlike statistical methods, which simply choose between or generate answers, my techniques can be inspected to understand *why* a particular answer was generated. Though recent work on statistical methods promises greater explainability (Branting, et al., 2019), there is still no substitute for examining the internals of a system. One can inspect the analogy to determine what entities and expressions were placed into alignment, to determine whether the algorithm stumbled into a correct answer through blind luck or derived it by properly placing the case at bar into alignment with the prior case. And this inspectability and explainability are not useful only for research debugging, but to constructing a functioning system: a legal reasoning system that can explain itself should be more trustworthy and therefore more useful than one that cannot.

Most importantly, the RRLG rule-learning algorithm performs as well as or better than all the other approaches tested using Precision@1 or when using a strict truth test. While the improved performance of systems using a partial truth test suggests that there is more information the rule-learning algorithm could be capturing, this nonetheless represents an important result: a system that compares legal precedents to identify shared legal principles; converts those principles into rules; and applies those rules to a case at bar.

D. FUTURE WORK

Several areas of future work have already been mentioned. One is to generate cleaner generalizations, such that these methods can reliably use only the strict struth test. Another concerns retrieval, that is, ensuring that the best case is retrieved the first time around. This is not only to avoid Precision@N testing, but because in the real world a system will of course not have ground truth to compare its answers to. Also, new large-language-models are constantly being released, and it would be instructive to evaluate the performance of, for example, ChatGPT on our dataset instead of just GPT-J.

The Illinois Intentional Tort Qualitative Dataset also includes a little over a dozen self-defense cases, and I want to investigate how to apply these approaches to those cases. Self-defense presents an interesting challenge, in that it requires at least two reasoning steps: the first to determine that a party *could* be liable for tortious action, and a second to determine that the action is excused through an affirmative defense. I also wish to investigate whether performance remains when testing cases against true precedents, i.e., only providing the system with cases with which to reason that are prior to the case at bar. Finally, we want to leverage the ontology and knowledge base to rerepresent cases and reason about novel situations, to allow us bring the full power of analogical reasoning to bear on the legal domain.

CHAPTER 7: LEGAL ARGUMENTATION.

This chapter describes a system that uses the rules generated by the system (described in the previous chapter) for legal argumentation, i.e., a lawyer's argument for an outcome favoring a particular side. It reports the results of a pilot experiment and describes future areas of inquiry.

Precedential legal reasoning of the kind that is the subject of this thesis involves applying legal precedents to the fact of a new case to determine how the new case should come out. When performed by judges, this analysis involves seeking the best, clearest possible understanding of both the facts and the law, in order to arrive at the "right" decision in a given case. But this kind of reasoning is essentially the same reasoning used by lawyers arguing for one side or another, with the difference being that the lawyers have some desired outcome in mind already when the reasoning process begins (namely, "my client wins"). Instead of taking the most objective possible view of the facts and of the law, the lawyer leverages ambiguity in both to present both the facts and the law such that the law (as she presents is) as applied to the facts (as she presents them) will lead to an outcome favorable to her client.

This observation is not a criticism of lawyers, simply a description of their role, and the fact that lawyers work in areas of ambiguity does not imply any misrepresentation of facts or law. Laws overlap and contradict each other, and lawyers can reasonably argue that the governing law between two overlapping laws should be the one that happens to lead to the best outcome for their client. Sometimes the law is genuinely ambiguous or a case presents a new configuration of facts that requires new lawmaking, and lawyers can craft arguments about what the law should (or must) be. But most often the law is perfectly clear, and lawyers are instead arguing and building a narrative around what facts occurred. Indeed, trials are fundamentally fact-finding exercises designed to arrive at the truth of some set of events, with the law governing those events clear, and the outcome therefore set once the facts are revealed. And the facts underlying the events are the subject of a trial can be genuinely ambiguous: at what point did the aggression in the voice and demeanor of some participant in an argument rise to the level of imminently threatening, and therefore justify the use of force? Where *is* the property line between these two plots, when multiple deeds going back decades or centuries describe each in different ways? Was it "reasonable" to slow down only so much when going into the blind curve? While lawyers are bound by codes of ethics and are not allowed to dissemble to the court, the lawyer's job is to faithfully represent the best interests of their client, and to work in the grey areas that the facts and the law leave undetermined.

Lawyers engaging in precedential reasoning must use much the same tools that judges do. Much as judges do, lawyers must determine what law the precedents governing some legal doctrine have laid down. The difference is that where judges do not seek to expand the grey areas left undetermined by laws (and where a case strikes a grey area, seek to do what is "right"), lawyers make use of those grey areas to argue in favor of outcomes beneficial to their clients. They might argue that a grey area does not exist, when settled law in similar (but saliently different) cases might lead to an outcome beneficial to their client, or they might seek to expand a grey area to argue that some outcome is not actually required by prior law. And while both judges and lawyers present a specific set of facts that lead to a particular outcome, for a judge the outcome follows from the facts,[48] whereas lawyers must present the facts in such a way as to lead to a favorable outcome for their client.

[48] Sometimes a judge might only be able to determine what facts occurred after receiving the verdict from the jury and seeing whose testimony they must have credited. And judicial opinions are written persuasively, with the goal of convincing the parties that justice was done and convincing a higher court not to overturn the decision. This task of determining what counts as a fact based on the outcome given by the jury is nonetheless fundamentally different

Lawyers and judges therefore begin from much the same perspective: examining precedent cases to determine what law governs some claim. They also examine the facts in a case in light of that law, to determine what outcome the law might demand. The difference is that lawyers are looking for opportunities to cast the law or the facts in a different light, to achieve a particular outcome. Given the similarities between these two reasoning processes, it should be possible to adapt a system developed for legal judgment into one developed for legal argument.

This chapter presents an initial implementation of a legal argumentation system adapted from the rule-learning system presented in Chapter Analogical Generalization, Reasoning, and Rule Learning for Legal Reasoning about Common Law Torts.. It presents a pilot assessment of the system. It then describes how a more thorough evaluation of the system would proceed, as well as how the other legal reasoning systems presented in Chapter Analogical Generalization, Reasoning, and Rule Learning for Legal Reasoning about Common Law Torts. could be adapted for legal argument rather than pure legal analysis. Due to time constraints these adaptations—and further evaluation—were impossible in the context of this thesis.

A. RULE REASONING FOR LAWYERS

I adapted the rule learning and reasoning system from Chapter Analogical Generalization, Reasoning, and Rule Learning for Legal Reasoning about Common Law Torts. to make legal arguments like those a lawyer would make. For each case, this system generated arguments for both sides of the issue. Because every case had a “right” answer that included the reasoning a judge would (presumably) use to govern that case, one side always had an argument in the form of the reasoning that the judge would presumably find persuasive. That is, to the extent the system could

from a lawyer determining what facts to present based on the desired outcome, because an impartial judge is not supposed to have a desired outcome in the case.

reason its way to the “correct” solution, it had also generated an argument for use by the lawyer seeking to achieve that particular outcome. For a positive case, the reasoning that justifies the legal conclusion—e.g., that the defendant trespassed—is also the argument in favor of the claimant—in that instance, the plaintiff. For a negative case, the reasoning that *fails* to justify a legal conclusion is the argument in favor of the side defending against the claim. The trickier question is, what to do for the other side in each of those instances?

For simplicity of discussion, let us assume that the side bringing a claim is always the plaintiff, and the side defending against it is always the defendant. (This is not always true: the dataset, for example, contains several cases where the defendant is accusing the plaintiff of having been trespassing, to avoid liability for injuries the plaintiff received on the defendant’s property.) In a positive case, the plaintiff claims that the defendant behaved tortiously, and they are correct; in a negative case, the plaintiff is incorrect. In a positive case, the plaintiff’s argument is the same analysis generated by the rule-learning system in Chapter Analogical Generalization, Reasoning, and Rule Learning for Legal Reasoning about Common Law Torts.; in a negative case, it’s the defendant’s argument that is the same. (Actually, a defendant’s argument in a negative case is that whatever the plaintiff is arguing is false – a subtle difference that is relevant to the evaluation, see below.) In a positive case, the defendant needs to argue that somehow the reasoning the plaintiff used is false, and in a negative case, the plaintiff needs to argue that, in fact, the facts necessary to draw a legal conclusion are present.

Start with the defendant in a positive case. He forms his argument by attacking the information the plaintiff used. This can be done by attacking the factual premises, or by attacking the legal rule itself (that is, claiming that the rule the plaintiff used is the wrong rule). My system does just that: it gathers the rule the plaintiff used along with the antecedent facts that the rule used to fire, and claims that one of those statements is false. Ideally the system would be able to preferentially select amongst these premises to attack: it would have some way of knowing whether a rule or a claimed fact stood on firmer or weaker ground, and attack the weakest premise. Even better, it would generate an argument as to *why* the premise should be attacked. But as an initial-cut system, it simply takes one of the premises at random and posits that it is not true. This might be a weak legal argument, but it is an argument. This approach is presented in Figure 7.1.

What of the plaintiff in a negative case? She forms her argument by identifying which facts are missing from her claim, i.e., which facts *would* have made the rule fire. This is trickier than a defendant's job in a positive case, because rather than grabbing a premise and saying "that is not true," the system has to identify what information is missing. Fortunately, FIRE has abduction as a basic reasoning capability. Abduction allows a system to assume missing information in service of inferring some consequent, and it can be used to essentially generate best-guess explanations for events. For example, you might have a rule that says "If it is raining, then the ground will be wet." When presented with wet ground, this rule can be fired backwards to abduce the fact that it was raining. We do not know that it was raining, we abduce it based on the presence of the wet ground.

It might seem perverse to say, "we do not know the defendant was on the plaintiff's property, we abduce it based on the fact that the defendant trespassed," but that is functionally

Figure 7.1: Arguing as a Defendant in a Case the Plaintiff Should Win.

Given positive case *c*, rule set *R*:
Defense(*c*, *R*):
1. **For** rule $r \in R$:
    2. **If** solvesCase(*r*, *c*): [Rule *r* can solve case *c*; attack its premises]
        3. *guiltyFacts* = {*f* for *f* in facts(*c*) where *f* participates in solvesCase(*r*, *c*)}
        4. *factsAndRule* = {*r*, *guiltyFacts*}
        5. *attackedFact* = randomElement(*factsAndRule*)
        6. **Return** "Not guilty because [*attackedFact*] is false."

what the plaintiff's argument generator does in a negative case. It queries the rules that *failed* to fire (and therefore conclude that the defendant committed some tort) with the rules' antecedents marked as abducible. The query will succeed by definition, installing in the system's working memory facts identified as having been abduced in service of that particular query. The system grabs those facts, and they become the plaintiff's argument. The system does this for every rule it has, and keeps as its argument the one that involves positing the fewest number of abduced statements. The fewer abduced statements being hypothesized, the closer the plaintiff was to having succeeded with a given rule. This approach is presented in Figure 7.2.

Again, in the ideal case the plaintiff would have some way of arguing *why* those facts are true. The system here is not generating the kind of argument one would want a lawyer to produce, because there is no justification of the particular facts being abduced. The system is instead simply identifying those facts as being missing, pointing the way for a human lawyer to generate the necessary argument justifying their presence.

For this reasoning system I did not use every rule generated by the rule-learning system. Instead, after creating the rules, I tested them against the very cases from which the rules were

Figure 7.2: Arguing as a Plaintiff in a Case the Defendant Should Win.

Given negative case *c*, rule set *R*:
Plaintiff(*c*, *R*):
1. *fewestAbducedFacts* = nil; *targetRule* = nil; *postulates* = nil [bookkeeping variables]
2. *conc* = plaintiffsDesiredOutcome(*c*)
3. **For** rule $r \in R$:
    4. abduceFromFactsAndRules(*conc, c, r*)
    5. *abducedFacts* = {*f* for *f* in WM abduced by abduceFromFactsAndRules(*conc, c, r*)}
    6. **If** *fewestAbducedFacts* = nil **or** count(*abducedFacts*) < *fewestAbducedFacts*:
        7. *targetRule* = *r*
        8. *postulates* = *abducedFacts*
        9. *fewestAbducedFacts* = count(*abducedFacts*)
10. **Return** " Guilty: [*postulates*] are true, so under rule [*r*] the defendant is liable."

created (Figure 7.3). I had the system count the number of cases that each rule was able to solve, and use the rules that solved the most number of cases (I took the rules with the top two scores, which may be more than two rules, in case of a tie). I scored the rules in an attempt to weed out poorly formed generalizations: if some rule was never—or only rarely—useful for solving a case, that could suggest that the schema from which the rule was built did not accurately capture the relevant information. Of course, edge cases do occur in the law and lawyers must be able to rely on applicable rules no matter how rarely they apply, so I do not claim that this scoring is a theoretically necessary component of my system.

B. EXPERIMENTAL VALIDATION

I ran this legal argument system on all cases in Assault, Battery, and Trespass, and manually inspected the results. Not every case would result in an argument. When a rule fails to fire, that failure indicates that the rule supports the defendant's claim; the rule would then be attempted as a target of abduction to generate an argument for the plaintiff. But while abduction is able to hypothesize a relationship between entities not present in a case, the abductive query will fail if it involves abducing entities not present in a reasoning context (that is, an abductive query with a variable that cannot be resolved to a known entity will fail rather than create a skolem

Figure 7.3: Selecting Rules for the Argumentation Rule Set.

Given case set *C*, rules derived from those cases *R*:
SelectRules(*C*, *R*):
    1. *topRuleScore* = nil; *secondBest* = nil
    3. **For** rule $r \in R$:
        4. ruleScore(*r*) = 0
        5. **For** case $c \in C$:
            6. **If** solvesCase(*r*, *c*): ruleScore(*r*) += 1
        7. **If** *topRuleScore* = nil **or** ruleScore(*r*) > *topRuleScore*: *topRuleScore* = ruleScore(*r*)
        8. **Elif** *secondBest* = nil **or** ruleScore(*r*) > *secondBest*: *secondBest* = ruleScore(*r*)
    10. **Return** {*r* for *r* in *R* where ruleScore(*r*) >= *secondBest*}

variable to resolve it). If every rule fails *and* is unabducible, then the plaintiff is unable to make an argument in her own favor using my argumentation system.

For negative cases, the defendant's argument is that whatever argument the plaintiff is making—whatever facts the plaintiff is abducing—is false. That means that in my system, when a plaintiff is unable to make an argument, the defendant is unable to as well. Now, it would be trivially easy to have the system generate a default "the defendant wins" argument, and indeed I considered having it do just that. But I believe that evaluating the system's failure to generate an argument based on the actual facts and rules is a more effective assessment of the system's argumentation abilities than would be returning a pithy "the defendant wins by default" response.

The system was able to generate arguments in 58 cases (58%).[49] This is higher than RRLG's performance at *solving* cases (47%), with good reason: cases where RRLG solved the case improperly by firing a rule are ones where the argumentation system is able to generate an argument. On the other hand, the argumentation system can also be used to assess accuracy, meaning, I can keep track of whether the system was able to correctly fire rules on positive cases and fail to fire them on negative cases. Here, the system only generated 31% accuracy, much lower than RRLG's 47% accuracy. That indicates that the lower-ranked rules filtered out by the argumentation system were in fact useful for solving the cases from a reasoning perspective.

The argumentation system was able to generate arguments—meaning the rules either fired or were abducible—in 34 Trespass cases (79%, equal for positive and negative cases), 12 Assault cases (71%, arguing 75% of positive cases and 60% of negative cases), and 12 Battery cases (only 30%, arguing 33% of positive cases and 23% of negative cases). The arguments suggest that there

[49] When the system fails to generate an argument, it returns nothing. Example arguments are provided below.

will be significant advantages to investigating the ability to re-represent case facts. For example, several arguments made by plaintiffs in positive cases where the rule failed to fire include facts that actually occurred, but were not represented in the case the way the rule expected them. For example, the case *Marks v Custom Aluminum Prods* (framed positive) involves a boss hitting and damaging his employee's arm. CNLU faithfully translated those facts, but does not have in the case representation that damage to the plaintiff's arm damages the plaintiff, or that the hit is an attack on the plaintiff. It also represents the actor role relation as being doneBy, rather than performedBy. The abduced argument in favor of the plaintiff in that case was `(and (damages hit9326 plaintiff29182) (objectAttacked hit9326 plaintiff29182) (isa hit9326 AttackOnTangible))`.[50] With an effective rerepresentation system that could reason about the events involved, the rule would have fired in the original instance.

Examining the arguments generated is encouraging. Where rules failed to fire (but are abducible) in positive cases, the plaintiff's side consistently generates arguments that contain facts that would establish the claim and are not absurd (and again, many that would be found true with an effective rerepresentation system). The same is true where abducible rules failed to fire in negative cases: the system proposed, as arguments in negative assault cases, the firing a gun in celebration was a threat to a case's plaintiff; that imprisoning a plaintiff threatened them; that punching a plaintiff threatened them. That said, in the trespass cases the plaintiff's arguments in negative cases sometimes made little sense: for example, in a case where a newspaper's employee went to retrieve printing machines from a warehouse and injured himself when he fell down a hole looking for the bathroom (he was not trespassing), the system proposed as an argument `(and`

[50] To be useful in the real world, the arguments generated by these systems should be in English, not predicate logic. This challenge is a natural language understanding and generation problem, and outside of the scope of this thesis.

`(doneBy machine94342 plaintiff93749) (doneBy look95625 plaintiff93749)).` (Here, the plaintiff was the one being accused of trespassing, as a defense against liability.) The entity `machine94342` is in the case as an element of the `group-of-machine94342` that the plaintiff was sent to retrieve; it makes no sense to suggest that the machines were "doneBy" the plaintiff.

Examining defendants' arguments in cases where rules successfully fired reveals how critical it would be to have the system understand what facts are attackable. For every argument of the form `(not (doneBy destroy17908 defendant15339))` or `(not (to-Generic enter59803 property59889))`, there several of the form `(not (isa land98425 RealEstate))`, `(not (isa property44983 Property))`, or even `(not (isa die68193 Dying))`. These would be laughably obviously false arguments in a courtroom. The current system has no way of distinguishing between terrible arguments ("that land is not property"), possibly acceptable arguments ("you do not own that property"), and good arguments ("I was not on that property"). Of the 26 cases where the system was able to fire rules and generate arguments for the plaintiff, in 21 cases the defendant's argument involved attacking premise facts and in 5 cases it involved attacking the rule. Again, the system was randomly choosing what to attack so there is not much to read into those numbers except as a measure of the size of the rules; a more sensible premise-attacking system would be extremely beneficial.

The results of this pilot experiment suggest that the argumentation system as it currently exists will be much more effective for determining what facts are missing from a claimant's case, and directing their efforts towards those facts, than to generate arguments defending against otherwise successful claims. This system could also be useful to direct efforts in building a rerepresentation and commonsense reasoning system, by revealing what facts are not being found that a cursory inspection by a human reveals should be true.

This is only a pilot experiment and not done with the degree of rigor required for publication. At a minimum, other evaluators than myself should be examining and evaluating the arguments generated by the system. The rules themselves should also be examined by lawyers to determine their fidelity to actual legal rules. But these results nonetheless demonstrate that the same rule-learning system that can be used from the perspective of a judge to resolve cases can be used from the perspective of a lawyer to argue them.

C. FUTURE WORK IN ARGUMENTATION SYSTEMS

Besides conducting a more thorough evaluation of the initial argumentation system, I could also adapt the rule-learning system so that when someone defending a claim seems to lose under one rule, they could instead propose a different rule under which they would win. It should also be fairly straightforward to adapt my legal reasoning systems that use analogy to generate arguments rather than only objective analysis.

As with the rule-reasoning system, the argument for one side is always provided by the system's success or failure to infer that the accused party is, according to its application of precedents, liable. If PAPR or ARPG infers that a party has committed a tort, then that becomes the argument for the accuser; if it infers that a party has not, that becomes the argument for the accused. And as with RRLG, the arguments for the other side can be generated based on what does (or does not) participate in the first argument. But the available arguments are different from RRLG.

First, in RRLG an argument available to the defense is that the rule used is the wrong rule. In PAPR or ARPG, the system could instead propose a different precedent with which to reason. This would especially be useful in PAPR, which has access to both positive and negative cases: every time a plaintiff wins by analogy to a positive case, the defendant can counter with an analogy to a negative case, and vice versa. The challenge there would be in ensuring that the analogies generated are actually accurate. On the other hand, this would represent essentially the same mechanism used by CASI to learn (i.e., ensure a mapping generates a particular candidate inference), so there is every reason to think it would be an effective argument generator.

Arguing for a different precedent generally involves *distinguishing*, identifying *why* some precedent should not apply (and perhaps that a different one should). It is not enough to argue that a precedent generates the wrong outcome! One way to handle distinguishing is as HYPO does: identifying facts present in the current case and absent in the precedent (or present in the precedent and absent in the current case), and proposing that the difference is sufficient to make the precedent inapposite (Ashley K. D., 1999). Another way is to incorporate reasons into the rules (Horty, 2011). In Horty's reason-based model, a case outcome (and the rule that derives it) comes not from the rule itself, but from the *reasons* for the rule. If rules are grounded in the reasons that justify them, and the reasons can be given a preferential order, then rules can be adapted to new circumstances based on which competing reasons must be given precedence. Reasons have not been a part of the case representations in the datasets used by my systems; it would be interesting to investigate whether incorporating them as causal drivers of case outcomes would be effective.

As with RRLG, parties can also attack or hypothesize facts in cases. For ARPG attacking facts would be straightforward, because a party defending against a claim would simply identify a salient fact participating in the claim and attack it, as in RRLG. It is even possible that attacking the right fact would lead to the mapping that produced the conclusion no longer being properly generated. Similarly, in a mapping that produced a particular conclusion with extra candidate inferences (i.e., a negative conclusion for ARPG), a putative plaintiff could posit the existence of the missing facts in the case.

Attacking or proposing facts would be trickier for PAPR, because abducing a fact that would make an analogy effective may be less straightforward than abducing a fact demanded by a rule. If reasoning directly from a past case rather than from a schema, it would not be obvious how to identify the facts that, if present, would nudge a mapping to generate the “right” conclusion. This concern would be less salient if the cases contained the kinds of higher-order structure SME is most effective at leveraging, but would remain a concern if using cases from the Illinois Intentional Tort Qualitative Dataset. Attacking facts would be simpler: the system could identify expressions that participate in the mapping and regenerate the mapping were that expression removed. The set of expressions that, if removed, would lead the mapping not to generate the conclusion statement become the set of facts that can form the basis of arguments on the attack.

To make all these systems effective and reasonably performing, there must be some way to determine what is a fact that is reasonable to attack and what is not. One possible heuristic is only to attack facts that directly involve their “client,” although this may be needlessly limiting. But figuring this out will be critical to making these systems effective defenders. Otherwise they risk being like the borrower in the old joke, who when told, “Hey, you returned my pot to me broken,” responds, “First, it was in perfect condition when I returned it; second, it was broken when you gave it to me; and third, I never borrowed your damn pot!”

CHAPTER 8: LIMITATIONS AND FUTURE WORK.

This chapter discusses limitations inherent to using certain kinds of computational legal systems, and of my own systems. It discusses limitations related to requiring specified case facts, and to the challenges posed by open-textured terms.

First, the systems herein are only tested on simple Tort cases, and have not yet been tested on the self-defense affirmative defense cases in the Illinois Intentional Tort Qualitative Dataset. Affirmative defenses require a multi-step reasoning process, first to recognize that some claim should be valid, and then to explain why the behavior is excused or justified. The Dataset should also be extended to include the contract cases that were collected in the initial dataset stage, and the systems presented here should be tested on those cases.

Another clear limitation is in the quality of the generalizations learned. While my experiments demonstrate that the system is learning and applying decent-to-good representations of the legal principles governing certain doctrines, the generalizations do not always include all legally-relevant facts, and continue to include some legally-irrelevant ones. Legally irrelevant facts might be no problem for a system reasoning entirely by analogy (like PAPR), but it is a problem for a system like RRLG that depends on converting generalizations into Horn clauses. Furthermore, all the cases in the dataset were selected for being similar to at least one other case in the dataset; the sheer number of outliers suggests that there is more information for CASI (or SAGE) to be capturing. Getting more cases to generalize and ensuring the quality of the generalizations is a significant area for future work. Improving the generalizations may also reveal that the theoretical assumption that legally-irrelevant facts will fall away from legal generalizations is wrong. If so, I will have to investigate how to identify and remove legally irrelevant facts from generalizations before converting them into Horn clauses (or using them for ARPG).

The limitation regarding generalizations is an immediate and practical one. However, the approach described in this thesis also has several theoretical limitations.

A. WHERE THE SIDEWALK ENDS

The point has been made repeatedly and about a huge variety of AI systems that a machine learning system that learns to reason only from the past will be stuck with the patterns and reasoning encoded in its training data (O'Neil, 2016). While different ML systems might be better or worse able to adapt their learned reasoning to new stimuli, they are unable to make a normative judgment regarding when a decision rule *should* change, let alone craft a new rule to supplant the old one. Even in their theoretically most effective and most ideal forms, the systems described in this thesis will learn and apply historical rules, and be unable to craft new ones from first principles.

Still, there are steps that can be taken to make such systems more flexible in the face of new information. Being able to rerepresent case facts at higher (or lower) levels of abstraction would allow the systems that reason by analogy to recognize when new case facts that seem to present new situations ought nonetheless to be governed by old rules that otherwise do not seem to apply to them. If this rerepresentation focused on the learned schemas or precedents, it could lead the system to understand what it had learned in a different light, adapting its precedents in a way that would make them applicable to the new case. If the rerepresentation focused on the new case, it could lead the system to recognize that what appeared to be a new set of circumstances actually broadly fit within the old set of circumstances. And rerepresenting both the precedent and the case at bar simultaneously would allow the system to search for systematic representations that make the old and the new cases congruent. To the extent that such rerepresentation is possible, it could also help develop better generalizations, by allowing facts that should generalize together (e.g., the predicates doneBy and performedBy) to do so. Better generalizations should lead to better performance for ARPG and RRLG.

The challenge with rerepresentation is how to know when to do it. In the Cyc ontology, just about every entity can be rerepresented as a Thing, because most collections eventually inherit from Thing. But rerepresenting cases as a Thing doing a Thing to another Thing's Thing does not seem like it would usefully capture information about events. Nor is the same level of generality always appropriate: when learning about rules for bringing pets on board an airplane, perhaps rerepresenting cats and dogs as Animals might make sense, whereas when learning about what kinds of games to play with them, keeping them separate will avoid the significant frustration of trying to get one's cat to Fetch.

One way to move forward on this problem may be to identify the predicates and entities that are often good candidates for the result of rerepresentation: objectActedOn, doneBy, etc. One would not always want to rerepresent these predicates because they can make very different situations seem quite similar (by, again, making every situation one in which a thing did something to something else). This is already a problem with the outputs of CNLU, where objectedActedOn and performedBy are generated by a huge number of verb semtranses. But it could help during the matching phase, to help the system recognize when two dissimilar cases might be closer together than it seems. It also suffers a knowledge-acquisition bottleneck problem if the list must be hand-generated.

Another option is to use minimal ascension, which allows SME to rerepresent predicates at a higher level of the ontology. Minimal ascension can make use of a rerepresentational depth limit that the system is willing to walk, just to try rerepresenting things. There are two possible problems with this: first, it is potentially inefficient, since not everything is a good target for rerepresentation. The second problem is that the ontology is not uniformly densely populated. Some subjects feature highly granular distinctions at different levels of abstraction, while others make large leaps in conceptual abstraction. This is often due simply to where research attention was focused on developing the hierarchy. As a result, it may be hard to find a sweet spot number of abstraction steps where concepts from hierarchically dense parts of the ontology actually get rerepresented in meaningful ways, but concepts from the sparse parts of the ontology don't all just turn into Things in a few short leaps up the hierarchy.

Second order analogies could also help adapt existing doctrines to deal with new circumstances. If some new set of circumstances cannot be rerepresented to fit within existing doctrines, then the new facts can be analogized to different fact patterns in known cases to see which doctrine ought to govern them. That is, if I have precedents A and B that lead to different outcomes, and some new fact pattern C comes in that cannot be rerepresented as either A or B, then the individual components of C can be analogized to the components of A and B, to determine which one will be a more appropriate analog. This can be done through second order analogies: the new case C can be compared to both A and to B, and those comparisons can be compared to each other to determine which is a better fit.

Recognizing when a doctrine must be overruled is harder still. As noted in Chapter 3: there are compelling reasons to have the crafting of legal rules only be performed by humans. Even identifying when a new rule might be appropriate (without regard to what that new rule ought to be) is an extremely challenging task. Perhaps a system that is able to reason about moral rules and norms (Olson, 2021) could identify when applying the old rules would lead to a norm violation, and flag its decisions as ones where human review would be necessary.

B. SEEKING FACTS

There is a way in which vast swaths of AI & Law research, including my own, puts the cart before the horse for legal reasoning. As mentioned previously, trials are fundamentally fact-finding endeavors. The law is usually pretty clear, and the question is largely about figuring out what it is that happened, so the judge can determine what outcome the law requires given the facts that have come out at trial. Indeed, a huge number of cases are resolved at the summary judgment stage, where the parties have presented the judge with enough facts upon which they agree that the judge can determine what outcome is required, and a trial is not even necessary.

AI legal reasoning systems (including my own), on the other hand, generally start with a set of facts predetermined. This assumption – that the set of facts can be known – assumes away most of the actual legal process. Maybe that is fine, since I am not trying to build a system to replace the legal process: if the point is just to build a system that can do the work that a judge will do *given* a set of facts, or to advise someone of what their legal rights or liabilities might be *given* a set of facts, then needing to insert those sets of facts could be a reasonable prerequisite to building a system. The place where it gets tricky is when one is trying to build a system to generate legal arguments. There are certainly legal domains where the law is sufficiently muddled or contradictory, or where cases have come out in sufficiently contradictory ways, that one can make an argument that different *law* should govern some case (or, put differently, that a different precedent should govern the case). But in situations where the law is clear, lawyers need to be able to argue the facts themselves. If so, the best that a system that assumes facts as inputs might be able to do is what I did with my argumentation system: to point to which facts are missing or are unfortunately present as a way to direct the lawyer's work. The actual work of arguing – of creating a reasoning framework that is in support of or against the existence of certain facts – will be assumed away or left to future work.

Describing the role of the lawyer does not even address the biggest problem with requiring facts as inputs, which is the assumption that facts are knowable. We need to have trials for a reason: facts can be hard to discern. Often they come down to a judgment call by a jury, whether about a particular person is trustworthy and therefore has their account accepted, or whether some action was “reasonable” as the jury understands it. Needing to know the facts means giving away most of the game. It is one thing to say that a system that learns legal rules will do so from the facts described by the judge in precedent opinions, and another to say that the system will *apply* those rules to cases with known facts. The facts are only known for sure once the trial is over and a verdict has been rendered. The facts are released in an opinion along with the reasoning and the law applied to them.

It may be possible to make some progress on adapting systems like the ones proposed here to deal with ambiguities in facts. One way might come, again, through rerepresentation. Everyone might agree that one person raised their hand, but might disagree over whether it was in a threatening gesture versus to ask for the opportunity to speak during an argument. Having a rerepresentation system that could move between different representations of the same action could allow a legal argument system to work from the same set of facts to create different arguments and interpretations of those facts. Another method might be through the development of novel case constructors: a case could be represented initially not as a set of grounded facts, but as a set of all *possible* facts for that case, and case constructors could generate interpretations of that case based on consistent sets of those potential facts. All of this is left for future work.

C. OPEN TEXTURE

The law features a bevy of open-textured terms, terms that are either purposely left vague by a statute drafter or court to recognize the limits of advance planning and to give later decision-makers leeway to adapt to the circumstances they face, or turn out to unavoidably be vague due to the inherent ambiguity of language (Hart H. , 1961). Open-textured terms were an area of focus in the early days of AI & Law (e.g., (Rissland & Skalak, 1988; Sanders, 1995; Bench-Capon, 1993)), but focus has fallen away from them more recently. While open-textured terms are most often associated with statutes, they are also present in common law legal doctrines, most prominently exemplified in the word “reasonable.”

The law of Negligence says that person A is responsible for an injury to person B if person A owed a duty of care to person B that they breached, and that the breach caused person B’s injury. While there are many specific duties of care (e.g., the care a parent owes their child), there is also a generalized duty of care to behave “reasonably”. What defines reasonable behavior? Ask a jury: whether behavior is reasonable is exactly the kind of “fact” that is adduced at trial, but is also in many circumstances inherently a judgment call.

I do not have a solution to propose to this problem. Perhaps a sufficiently large deep learning system could be trained to output judgments that track with people’s general intuitions about whether something is reasonable or not. But I do not think it would be a good idea to entrust such a system with the power to make any sorts of legal judgments, and the finding of whether behavior is reasonable is indeed a legal judgment. For one thing, unless we know exactly how the system is going to work -- and we generally will not -- we will be unable to predict strange edge cases where the system outputs an answer everyone would agree is wrong. For another, while being able to know in advance whether behavior will be found illegal can be useful – and indeed at least one commenter has argued that precisely knowing what would be illegal would be a boon and argues in favor of such systems, (Genesereth, 2015) – it strikes me as dangerous to be able to precisely tailor inputs to a system to get your behavior to fall *just on this side* of legal. If one knows exactly how some system will perceive their behavior, then they can game that system (Pasquale, 2015; O'Neil, 2016). Indeed it is exactly the ability to adapt to situations that makes open-textured terms useful; having a computational system that functionally does away with them will have unintended consequences.

My systems deal with open-textured terms essentially by ignoring them: to the extent any are implicated in the legal doctrines I model, my systems learn the boundaries of those open-textured terms directly from the training cases. And maybe that is enough of a solution for certain doctrines, although it is inherently backwards looking. On the other hand, without the ability to discern whether behavior is reasonable or not, I do not believe that such systems would be able to make much headway on cases in Negligence.

CONCLUSION

This thesis has presented three analogy-based models of precedential legal learning and reasoning. It introduced a new dataset – the Illinois Intentional Tort Qualitative Dataset – of Illinois Tort cases (in Trespass, Assault, Battery, and Self-Defense), the only currently publicly-available dataset of such cases represented in predicate logic. It described and evaluated a modification to the SAGE analogical generalization system that replaces SME's similarity score to control analogical generalizations when cases are dissimilar and feature low structural similarity scores. This modified algorithm not only supports learning legal rules from such cases, but will enable future researchers facing the same problem to learn useful concepts from their datasets. It introduced and evaluated three legal precedential reasoning systems: one that reasons by pure analogy to positive and negative cases (and schemas thereof); one that reasons *about* analogies made to schemas of positive cases; and one that converts those schemas into Horn clauses and reasons using legal rules. These algorithms present a new approach for analogical legal reasoning and provide an account of how legal rules might be discerned from a body of legal precedent. The thesis also described an implemented adaptation of the legal rule-reasoning system that would generate legal arguments, not only legal analysis, and described how the other reasoning systems could be adapted for argumentation. Finally, the thesis explored generalized principles and limitations concerning AI modeling of legal reasoning, and when doing so is appropriate.

It is an exciting time to be working at the intersection of AI and Law; it is also a scary time. AI is pervading ever more corners of modern human existence. AI is flying drones, approving mortgages, sifting resumes, connecting friends and arranging loves, driving cars, setting ticket prices, choosing what news you consume, creating new thoroughfares in suburban neighborhoods, placing ads, picking your music, the list is endless. Many, but not all, of the items just enumerated are legally regulated and carry with them the possibility of legal repercussions if improperly performed. The law is also changing, along with peoples' relationships with the law and the legal structures around them. AI technologies – for example, that direct police to specific areas, distribute government resources, identify potential criminal behavior, and more – are involved in many of these changes, sometimes in ways that are visible to the regulated population, but other times not.

The thesis that I have written will, of course, do nothing to address any of these issues. But I write it with the hopes for a future that will see people having more power over their own lives, and a greater understanding of the structures that drive the world around them, and I hope that those motivations have come through in my writing. Not every AI system needs to be explainable, but the ones that make important decisions that affect peoples' lives should be. Though it has been less of a focus of this thesis, I believe the law has a much stronger role to play in regulating the ways in which AIs are used in the name of profit and efficiency. I believe people should be able to understand the laws to which they are subject, and they should not need to be wealthy for the privilege of doing so. AI systems that can help people understand their rights and liabilities and help them vindicate those rights have the potential to be a great positive force in the world (give or take increasing America's litigiousness).

The future is upon us. Many of the technological developments of the speculative fiction books I read as a child have come to pass. Global warming is no longer a threat from the future. Our institutions are not as stable as we thought they were. A networked virtual reality world called the Metaverse is here, but instead of being open-access and distributed, it is a private product owned by one of the world's largest corporations. But humanity is not powerless, condemned simply to watch a slow decay. The tools that have been the instrument of this progression can also be those of the progress away from it. The law has long been the infrastructure of society; AI is becoming infrastructure in our daily lives. These systems are not broken: they just need our support in order to support us all in return.

REFERENCES

Al-Abdulkarim, L., Atkinson, K., Bench-Capon, T., Whittle, S, Williams, R., & Wolfenden, C. (2019). Noise Induced Hearing Loss: Building an Application Using the ANGELIC Methodology. *Argument and Computation, 10*(1), 5-22.

Alberdi, E., Ayton, P., Povyakalo, A., & Strigini, L. (2005). Automation bias and system design: A case study in a medical application. *People and Systems - Who Are We Designing For*.

Aleven, V., & Ashley, K. D. (1997). Evaluating a learning environment for case-based argumentation skills. *ICAIL '97: Proceedings of the 6th international conference on Artificial intelligence and law,* (pp. 170-179).

Alexander, L., & Sherwin, E. (2008). *Demystifying Legal Reasoning.* Cambridge University Press.

Allen, J. F. (1994). *Natural Language Understanding* (2d ed.). Redwood City, CA.

Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016, May 23). Machine Bias. *ProPublica*.

Ashley, K. D. (1989). Toward a Computational Theory of Arguing with Precedents. *Proceedings of the Second International Conference on Artificial Intelligence and Law* (pp. 93-102). ACM.

Ashley, K. D. (1991). *Modeling Legal Arguments: Reasoning with Cases and Hypotheticals.* MIT Press.

Ashley, K. D. (1999). Designing Electronic Casebooks that Talk Back: the CATO Program. *Jurimetrics, 40*, 275.

Ashley, K. D. (2009). Ontological Requirements for Analogical, Teleological, and Hypothetical Legal Reasoning. *Proceedings of the 12th International Conference on Artificial Intelligence and Law.* ACM.

Ashley, K. D., & Bruninghaus, S. (2006). Progress in Textual Case-Based Reasoning: Predicting the Outcome of Legal Cases from Text. *Proceedings of the Twenty-First AAAI Conference on Artificial Intelligence.* Boston, MA.

Atkinson, K., & Bench-Capon, T. (2019). Reasoning with Legal Cases: Analogy or Rule Application? *Proceedings of the 17th International Conference on AI & Law*, (pp. 12-21).

Barbella, D., & Forbus, K. D. (2015). Exploiting Connectivity for Case Constructing in Learning by Reading. *Proceedings of the Third Annual Conference on Advances in Cognitive Systems.*

Bench-Capon, T. (1993). Neural networks and open texture. *Proceedings of the 4th international conference on Artificial intelligence and law.*

Bench-Capon, T., & Gordon, T. F. (2009). Isomorphism and Argumentation. *Proceedings of the 12th International Conference on Artificial Intelligence and Law.* Barcelona, Spain.

Bengio, Y. (2012). Deep Learning of Representations for Unsupervised and Transfer Learning. *Proceedings of ICML Workshop on Unsupervised and Transfer Learning*, (pp. 17-36).

Blass, J. A. (2022). Observing the Effects of Automating the Judicial System with Behavioral Equivalence. *South Carolina Law Review*.

Blass, J. A., & Forbus, K. D. (2015). Moral Decision-Making by Analogy: Generalizations vs. Exemplars. *Proceedings of the 29th AAAI Conference on Artificial Intelligence.* Austin, TX.

Bowdle, B., & Gentner, D. (2005). The Career of Metaphor. *Psychological Review, 112*, 193-216.

Branting, L. K. (1991). Building Explanations from Rules and Structured Cases. *International Journal of Man-Machine Studies*, 797-837.

Branting, L. K. (2003). A Reduction-Graph Model of Precedent in Legal Analysis. *Artificial Intelligence, 150*, 59-95.

Branting, L. K., Weiss, B., Brown, B., Pfeifer, C., Cakraborty, A., Ferro, L., & Yeh, A. (2019). Semi-Supervised Methods for Explainable Legal Prediction. *Proceedings of the 17th International Conference on AI & Law*, (pp. 22-31).

Brewer, S. (1996). Exemplary Reasoning: Semantics, Pragmatics, and the Rational Force of Legal Argument by Analogy. *Harvard Law Review, 109*(5), 923-1028.

Bruninghaus, S., & Ashley, K. D. (2003). Predicting Outcomes of Case Based Legal Arguments. *Proceedings of 9th International Conference on AI & Law.* Edinburgh, Scotland.

Bussone, A., Stumpf, S., & O'Sullivan, D. (2015). The Role of Explanations on Trust and Reliance in Clinical Decision Support Systems. *International Conference on Healthcare Informatics (ICHI)*, (pp. 160-169).

Calo, R., & Citron, D. (2021). The Automated Administrative State: A Crisis of Legitimacy. *Emory Law Journal, 4*, 797-845.

Catrambone, R., & Holyoak, K. J. (1989). Overcoming Contextual Limitations on Problem-Solving Transfer. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*(6), 1147-1156.

Chalkidis, I., Fergadiotis, M., Malakasiotis, P., Aletras, N., & Androutsopoulos, I. (2020). *LEGAL-BERT: The Muppets Straight Out of Law School.* arXiv preprint arXiv:2010.02559.

Chen, B., Stremitzer, A., & Tobia, K. (2022). Having Your Day in Robot Court. *Harvard Journal of Law & Technology, 36*.

Chernova, S., & Thomaz, A. (2014). Robot learning from human teachers. *Synthesis lectures on artificial intelligence and machine learning, 8(3)*, 1-121.

Christie, S., & Gentner, D. (2010). Where Hypotheses Come From: Learning New Relations by Structural Alignment. *Journal of Cognition and Development, 11(3)*, 356-373.

Citron, D. K. (2008). Technological due process. *Washington University Law Review, 85*, 1249-1313.

Deghahni, M., Tomai, E., Forbus, K., & Klenk, M. (2008). An Integrated Reasoning Approach to Moral Decision-Making. *Proceedings of the 23d AAAI Conference on Artificial Intelligence.* Chicago, IL.

Dehghani, M., Sachdeva, S., Ekhtiari, H., Gentner, D., & Forbus, K. (2009). The Role of Cultural Narratives in Moral Decision Making. *Proceedings of the 31st Annual Conference of the Cognitive Science Society*. Washington, D.C.

Dworkin, R. (1978). *Taking Rights Seriously*. Cambridge, Massachusetts: Harvard University Press.

Dworkin, R. (1986). *Law's Empire.* Cambridge, Massachusetts: Harvard University Press.

Fallon, R. (2005). Legitimacy and the Constitution. *Harvard Law Review*, 1787.

Felleisen, M. e. (2009). *Semantics Engineering with PLT Redex.*

Fitzgerald, T., Short, E., Goel, A., & Thomaz, A. (2019). Human-Guided Trajectory Adaptation for Tool Transfer. *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 1350-1358.

Floridi, L., & Chiriatti, M. (2020). GPT-3: Its Nature, Scope, Limits, and Consequences. *Minds and Machines, 30*(4), 681-694.

Forbus, K. D., & Hinrichs, T. (2017). Analogy and Relational Representations in the Companion Cogntiive Architecture. *AI Magazine*.

Forbus, K. D., Ferguson, R. W., Lovett, A., & Gentner, D. (2017). Extending SME to handle large-scale cognitive modeling. *Cognitive Science, 41*(5), 1152-1201.

Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science, 19*(2), 141-205.

Forbus, K. D., Hinrichs, T., Crouse, M., & Blass, J. A. (2020). Analogies versus Rules in Cogntiive Architecture. *Proceedings of Advances in Cognitive Systems, 2020.*

Gardner, J. (2015). The many faces of the reasonable person. *Law Quarterly Review, 131*(1), 563-584.

Genesereth, M. (2015). *Computational Law: the Cop in the Backseat.* Palo Alto, CA: CodeX: The Center for Legal Informatics Stanford University.

Gentner, D. (1983). Structure-mapping: A Theoretical Framework for Analogy. *Cognitive Science*, *7*(2), 155-170.

Gentner, D., & Landers, R. (1985). Analogical Reminding: A Good Match is Hard to Find. *Proceedings of the International Conference on Systems, Man, & Cybernetics*. Tucson, AZ.

Gentner, D., & Markman, A. B. (1997). Structure Mapping in Analogy and Similarity. *American Psychologist, 52*, 45-56.

Gentner, D., & Smith, L. A. (2013). Analogical Learning and Reasoning. In D. R. (Ed.), *Oxford Library of Psychology* (pp. 668-681). New York, NY: Oxford University Press.

Gentner, D., Loewenstein, J., Thompson, L., & Forbus, K. D. (2009). Reviving Inert Knowledge: Analogical Abstraction Supports Relational Retrieval of Past Events. *Cognitive Science*, *33*(8), 1343-1382.

Gentner, D., Rattermann, M., & Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive psychology*, *25*(4), 524-575.

Gick, M. L., & Holyoak, K. J. (1983). Schema Induction and Analogical Transfer. *Cognitive Psychology*, *15*, 1-38.

Gonen, H., & Goldberg, Y. (2019). Lipstick on a Pig: Debiasing Methods Cover up Systematic Gender Biases in Word Embeddings But do not Remove Them. *Proceedings of NAACL-HLT 2019* (pp. 609-614). Minneapolis, Minnesota: ACL.

Gowder, P. (2018). Transformative Legal Technology and the Rule of Law. *University of Toronto Law Journal*, *68*, 82-106.

Gowder, P. (2020). Is Legal Cognition Computational? (When Will DeepVehicle Replace Judge Hercules?). In e. Ryan Whalen, *Computational Legal Studies* (p. 215).

Grabmair, M., Ashley, K. D., Chen, R., Sureshkumar, P., Wang, C., Nyberg, E., & Walker, V. R. (2015). Introducing LUIMA: an Experiment in Legal Conceptual Retrieval of Vaccine Injury Decisions using a UIMA Type System and Tools. *Proceedings of the 15th International Conference on AI & Law.*

Grant, E., Peterson, J., & Griffiths, T. (2019). Learning deep taxonomic priors for concept learning from few positive examples. *Proceedings of the 41st Annual Conference of the Cognitive Science Society*, (pp. 1865-1870). Montreal, Canada.

Guerini, M., Pianesi, F., & Stock, O. (2015). Is It Morally Acceptable for a System to Lie to Persuade Me? *AAAI Workshop on Artificial Intelligence and Ethics*.

Gulwani, S., Hernandez-Orallo, J., Kitzelmann, E., Muggleton, S., Schmid, U., & Zorn, B. (2015). Inductive programming meets the real world. *Communications of the ACM, 58*(11), 90-99.

Halstead, D. T., & Forbus, K. D. (2005). Transforming between propositions and features: Bridging the gap. *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI)*, (pp. 777-782). Pittsburgh, PA.

Hart, H. (1961). *The Concept of Law.* Clarendon Law Series.

Hart, H. (1961). *The Concept of Law.* Oxford University Press.

Holyoak, K. J., & Thagard, P. (1997). The Analogical Mind. *American Psychologist, 52*(1), 35-44.

Horty, J. F. (2011). Reasons and Precedent. *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Law*. Pittsburgh, PA.

Hummel, J. E., & Holyoak, K. J. (2003). A Symbolic-Connectionist Theory of Relational Inference and Generalization. *Psychological Review, 110*(2), 220-264.

Jin, X., Li, M., & Ji, H. (2022). Event Schema Induction with Double Graph Autoencoders. *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, (pp. 2013-2025). Seattle, WA.

Kamp, H., & Reyle, U. (1993). *From Discourse to Logic: Introduction to Model-Theoretic Semantics of Natural Language.* Boston, MA: Kluwer Academic.

Kandaswamy, S., & Forbus, K. (2012). Modeling Learning of Relational Abstractions via Structural Alignment. *Proceedings of the 34th Annual Conference of the Cognitive Science Society (CogSci).* Sapporo, Japan.

Kannai, R., Schild, U. J., & Zeleznikow, J. (2014). There is More to Legal Reasoning with Analogies than Case Based Reasoning, but What? In N. Dershowitz, & E. Nissan, *Choueka Festschrift* (pp. 440-451). Springer-Verlag Berlin Heidelberg.

Krizhevsky, A., Sutskever, I., & Hinton, G. (2017). Imagenet classification with deep convolutional neural networks. *Communications of the ACM, 60*(6), 84-90`.

Kroll, J., Huey, J., Barocas, S., Felton, E., Reidenberg, J., Robinson, D., & Yu, H. (2017). Accountable Algorithms. *University of Pennsylvania Law Review, 165*, 633.

Kuehne, S., Forbus, K. D., Gentner, D., & Quinn, B. (2000). SEQL: Category Learning as Progressive Abstraction Using Structure Mapping. *Proceedings of the 22d Meeting of the Cognitive Science Society*, (p. 770).

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep Learning. *Nature*, *521*(7553), 436-444.

Lee, J., & See, K. (2004). Trust in Automation: Designing for Appropriate Reliance. *Human Factors, 46*(1), 50-80.

Legal Services Corporation. (2022). *The Justice Gap.*

Lenat, D. B. (1995). CYC: A Large-Scale Investment in Knowledge Infrastructure. *Communications of the ACM*, *38*(11), 33-38.

Levi, E. H. (1949). *An Introduction to Legal Reasoning.* Chicago, IL: University of Chicago Press.

Lind, E. A., & Tyler, T. R. (1988). *The Social Psychology of Procedural Justice.*

Lustig, C., Pine, K., Nardi, B., Irani, L., Lee, M., Nafus, D., & Sandvig, C. (2016). Algorithmic authority: the ethics, politics, and economics of algorithms that interpret, decide, and manage. *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, (pp. 1057-1062).

Marcus, G., & Davis, E. (2020). GPT-3, Bloviator: OpenAI’s language generator has no idea what it’s talking about. *MIT Technology Review*.

Markman, A. B., & Medin, D. L. (2002). Decision Making. In *Stevens' Handbook of Experimental Psychology*.

McCarty, L. T. (1997). Some Arguments About Legal Arguments. *ICAIL '97: Proceedings of the 6th international conference on Artificial intelligence and law*, (pp. 215–224).

McFate, C. J., Forbus, K. D., & Hinrichs, T. (2014). Using Narrative Function to Extract Qualitative Information from Natural Language Texts. *Proceedings of the 26th AAAI Conference on Artificial Intelligence.*

McLure, M. D., Friedman, S. E., & Forbus, K. D. (2015). Extending Analogical Generalization with Near-Misses. *Proceedings of the 29th AAAI Conference on Artificial Intelligence.* Austin, TX.

Morris, J. H. (1968). *Lambda-Calculus Models of Programming Languages (Ph.D. dissertation, MIT).* Retrieved from https://dspace.mit.edu/handle/1721.1/64850

Muggleton, S., & De Raedt, L. (1994). Inductive Logic Programming: Theory and Methods. *Journal of Logic Programming*, *19*(20), 629-679.

National Center for State Courts. (2015). *The Landscape of Civil Litigation in State Courts.* Retrieved from https://www.ncsc.org/__data/assets/pdf_file/0020/13376/civiljusticereport-2015.pdf

Novick, L. (1988). Analogical Transfer, Problem Similarity, and Expertise. *Journal of Experimental Psychology, 14*(3), 510-520.

Nozick, R. (1974). *Anarchy, State, and Utopia.*

Olson, T. &. (2021). Learning Norms via Natural Language Teachings. *Proceedings of the 9th Annual Conference on Advances in Cognitive Systems 2021.*

O'Neil, C. (2016). *Weaspons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy.* New York: Broadway Books.

Pasquale, F. (2015). *The Black Box Society: The Secret Algorithms That Control Money and Information.* Cambridge, MA: Harvard University Press.

Petrova, A., Armour, J., & Lukasiewicz, T. (2020). Extracting Outcomes from Appellate Decisions in US State Courts. *2020 JURIX Conference.*

Posner, R. A. (2006). Reasoning by Analogy (Book Review of Weinreb's Legal Reason). *Cornell Law Review, 91*, 761-774.

Prakken, H., Wyner, A., Bench-Capon, T., & Atikinson, K. (2013). A Formalization of Argumentation Schemes for Legal Case-Based Reasoning in ASPIC+. *Journal of Logic and Computation, 25*(5), 1141-1166.

Quinlan, J. R. (1990). Learning Logical Definitions from Relations. *Machine Learning, 5*(3), 239-266.

Rabold, J., Siebers, M., & Schmid, U. (2022). Generating contrastive explanations for inductive logic programming based on a near miss approach. *Machine Learning, 111*(5), 1799-1820.

Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., & Chen, M. (2022). *Hierarchical Text-Conditional Image Generation with CLIP Latents.* arXiv. Retrieved from https://arxiv.org/abs/2204.06125

Rawls, J. (1999). *A Theory of Justice* (Revised ed.).

Ribeiro, D. N., & Forbus, K. D. (2021). Combining Analogy with Language Models for Knowledge Extraction. *3d Conference on Automated Knowledge Base Construction.*

Rissland, E. L., & Friedman, M. T. (1995). Detecting Change in Legal Concepts. *Proceedings of the 5th International Conference on Artificial Intelligence and Law* (pp. 127-136). ACM.

Rissland, E. L., & Skalak, D. B. (1988). *Case-based reasoning in a rule-governed domain.* Amherst, MA: MASSACHUSETTS UNIV AMHERST DEPT OF COMPUTER AND INFORMATION SCIENCE.

Rissland, E. L., & Skalak, D. B. (1989). Interpreting statutory predicates. *Proceedings of the 2d International Conference on AI & Law.*

Rissland, E. L., Skalak, D. B., & Friedman, M. T. (1996). BankXX: Supporting Legal Arguments Through Heuristic Retrieval. *Artificial Intelligence and Law*.

Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*(4), 629.

Ross, B. H., & Kennedy, P. T. (1990). Generalizing from the Use of Earlier Examples in Problem Solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16*, 42-55.

Rottman, B. M., Gentner, D., & Goldwater, M. B. (2012). Causal Systems Categories: Differences in Novice and Expert Categorization of Causal Phenomena. *Cognitive Science, 36*(5), 919-932.

Ruppenhofer, J., Ellsworth, M., Schwarzer-Petruck, M., Johnson, C. R., & Scheffczyk, J. (2016). *FrameNet II: Extended theory and practice.* International Computer Science Institute.

Ruppenhofer, J., Ellsworth, M., Schwarzer-Petruck, M., Johnson, C. R., & Scheffczyk, J. (2016). *FrameNet II: Extended Theory and Practice.* International Computer Science Institute.

Sanders, K. E. (1995). *Chiron: planning in an open-textured domain.* Brown University.

Schauer, F., & Spellman, B. A. (2017). Analogy, Expertise, and Experience. *University of Chicago Law Review*, *82*, 249-268.

Steinberg, J. (2015). Demand Side Reform in the Poor People's Court. *Connecticut Law Review*, *47*(3), 741–805.

Strauss, D. A. (2005). Reply: Legitimacy and Obedience. *Harvard Law Review*.

Sunstein, C. R. (1993). On Analogical Reasoning. *Harvard Law Review*, *106*, 741-791.

Sunstein, C. R. (1996). *Legal Reasoning and Political Conflict.* Oxford University Press.

The Associated Press. (2000, April 25). *Jury Flips Coin to Decide Case*. Retrieved from AP News: https://apnews.com/article/2c7dd9f60106fe05b63e37b81ef7fe7b

The Associated Press. (2007, November 2). *VA. Judge Fired after Deciding Case with Coin Flip*. Retrieved from NBC News: https://www.nbcnews.com/id/wbna21600306

the New York Times. (1982, March 11). Judge Facing Charges in Coin-Flipping Case. *New York Times*, p. B3.

Tomai, E., & Forbus, K. D. (2009). EA NLU: Practical Language Understanding for Cognitive Modeling. *22d International FLAIRS Conference.*

Trench, M., & Minervino, R. A. (2015). The Role of Surface Similarity in Analogical Retrieval: Bridging the Gap Between the Naturalistic and the Experimental Traditions. *Cognitive Science, 39*(6), 1292-1319.

Verheij, B. (2017). Formalizing Arguments, Rules, and Cases. *Proceedings of the Sixteenth International Conference of Artificial Intelligence and Law*. London, UK.

Walton, D. (2010). Similarity, Precedent, and Argument from Analogy. *Artificial Intelligence and Law*.

Weinreb, L. L. (2016). *Legal Reason: the Use of Analogy in Legal Argument,* (2d ed.). Cambridge University Press.

Wyner, A. Z., Gough, F., Levy, F., Lynch, M., & Nazarenko, A. (2017). On Annotation of the Textual Contents of Scottish Legal Instruments. *JURIX 2017*, (pp. 101-106).

APPENDIX A: THE ILLINOIS INTENTIONAL TORT QUALITATIVE DATASET

Appendix A contains the first ten cases, alphabetically, of the Illinois Intentional Tort Qualitative Dataset. These cases are organized alphabetically rather than by doctrine since cases often include multiple doctrines. The complete dataset is not included because it would make this document over 1000 pages long. The entire dataset is available at the time of this writing at https://www.qrg.northwestern.edu/Resources/caselawcorpus.html; if you are reading this in the future and it is no longer available, I will have made efforts to preserve it online in a different place.

```
(in-microtheory CaseLawCorpusMt)

(isa Amos_v_State CaseLawCorpusCase)

(in-microtheory (CaseLawCorpusMtFn Amos_v_State))

(caseType Amos_v_State Assault)
(caseType Amos_v_State Battery)

(caseValence Amos_v_State Assault Positive)
(caseValence Amos_v_State Battery Positive)

(caseName Amos_v_State "Amos v. State")

(caseCourt Amos_v_State "Court of Claims of Illinois")

(caseYear Amos_v_State "2003")

(caseReporter Amos_v_State "55 Ill. Ct. Cl. 368")

(caseOriginalText Amos_v_State "On or about December 4, 1996 the plaintiff was
the owner of a two story residential building commonly known as 1541 Lincoln
Street located in the city of North Chicago, Illinois. the plaintiff, her family
and two pet dogs occupied the first floor unit. A tenant, who was the subject
of a drug investigation conducted by the Metropolitan Enforcement Group
(hereinafter MEG), a division of the Illinois State Police, occupied the
upstairs or second floor unit. On December 4, 1996, MEG obtained a court
authorized search warrant for the second floor residence. It is undisputed and
verified by the search warrant itself, as well as from elicited testimony, that
neither the plaintiff's first floor unit, nor anyone in her household was under
investigation or the target of the search warrant. On December 4, 1996 at
approximately 12:00 noon police officers arrived at the 1541 Lincoln Street
property in order to execute the search warrant for second floor unit. An arrest
of a male subject was subsequently effected for unlawful possession of cannabis,
```

which had no connection or bearing on the plaintiff. the plaintiff, a 40 year old manager for Goodwill Industries located off the Great Lakes Naval Base, testified that on December 4, 1996, she arrived home from work at approximately 12:15 p.m. At which time she released her pet dogs, Havoc a five year old Straffordshire Terrier and Chaos, a six month old American Pitbull Terrier, out into the backyard. After about five minutes she brought the dogs back into her first floor unit. the plaintiff, while changing clothes, heard smashing and banging sounds in her kitchen. She ran into the kitchen and observed the backdoor being slightly pushed opened. the plaintiff was able to close and lock the door while screaming for help. No one identified themselves as police officers at that time. She then ran into the bedroom, grabbed a cordless phone and ran for the front door. the plaintiff dialed 911 while running out of the house and onto an enclosed front porch. While on the enclosed front porch and talking to 911, the plaintiff opened the screen door which led directly to the outside and saw two men ascending the five steps that led up to the enclosed porch. They identified themselves as police officers, one of which was Officer George. the plaintiff informed 911 that the men at her door claimed to be police officers. 911 confirmed that they were. [Officer George] observed the plaintiff standing in the enclosed front porch and talking on the telephone. George identified himself as a police officer, walked up the steps and entered the enclosed porch. George testified that he listened to the plaintiff scream on the phone for several minutes apparently mad that the police had done a search warrant on the residence. Officers Bell, Parisi and Poulos joined George on the enclosed porch. While in the enclosed front porch, [Officers George and Poulos] heard and observed through a window pane of the front interior door that leads directly into the plaintiff's residence, two large dogs barking, scratching and jumping at the door. Poulos asked [the plaintiff] to go into the house and put the dogs away so [they] could speak. As the plaintiff opened the interior door in an attempt to put the dogs away, the dogs slipped passed her and ran out onto the enclosed porch. Even though the dogs got loose, the plaintiff managed quickly to grab and hold both dogs up against her chest. She and both dogs were about three to four feet away from Officer George as he stood upon a radiator. George then shot one of the dogs, while still in her arms. The dog lunged up into her body and then ran into the plaintiff's house.")

```
(caseConclusion Amos_v_State "Defendant commited an assault on the plaintiff when he shot the plaintiff's dog.")

(caseConclusion Amos_v_State "Defendant commited a battery on the plaintiff when he shot the plaintiff's dog while the plaintiff was holding the dog.")

(caseSimpleText Amos_v_State "The plaintiff owned a building. The plaintiff lived in the building. The plaintiff owned dogs. A tenant lived in an apartment in the building. The defendant is a police officer. The defendant had a search warrant for the tenant's apartment. The  plaintiff was not under investigation. The police arrested the tenant. The plaintiff heard loud noises in the yard. Then the plaintiff called 911. Then the plaintiff ran from the house to an enclosed porch. Then the defendant entered the porch. The defendant told the plaintiff he was a police officer. The plaintiff's dogs were in her apartment. The dogs  were barking. Then the dogs ran outside. Then the plaintiff quickly grabbed both dogs. The plaintiff held the dogs to her chest. The plaintiff and the dogs were not near the defendant. Then the defendant shot the dog while the plaintiff was holding it.")

(in-microtheory (rawLanguageOutputMtFn Amos_v_State))
(genlMt (rawLanguageOutputMtFn Amos_v_State) (CaseLawCorpusMtFn Amos_v_State))
```

(and (doneBy own48367 plaintiff48357) (isa building48405 RealEstate) (isa own48367 OwningSomething)
(objectActedOn own48367 building48405) (relationExistsInstance plaintiffs CourtCase plaintiff48357)
(isa building48405 Building) (isa plaintiff48357 SocialBeing))
(and (residence-Role live48468 building48405) (isa live48468 ResidingSomewhere) (residents-Role live48468 plaintiff48357)
(isa plaintiff48357 SocialBeing) (isa building48405 Building)
(relationExistsInstance plaintiffs CourtCase plaintiff48357))
(and (doneBy own48367 plaintiff48357) (isa own48367 OwningSomething) (objectActedOn own48367 dog48641)
(isa plaintiff48357 SocialBeing) (relationExistsInstance plaintiffs CourtCase plaintiff48357) (isa dog48641 Dog))
(and (objectFoundInLocation apartment48814 building48405) (residence-Role live48728 apartment48814)
(nameString apartment48814 building48405) (isa live48728 ResidingSomewhere)
(objectFoundInLocation tenant48681 apartment48814) (residents-Role live48728 tenant48681) (isa building48405 Building)
(isa apartment48814 ApartmentUnit) (isa tenant48681 Tenant))
(and (relationExistsInstance defendants CourtCase tenant48681) (isa tenant48681 SocialBeing)
(isa defendant48982 PoliceOfficer-Municipal))
(and (isa warrant49216 Examination-Investigation) (isa warrant49216 Warrant-Writ)
(possessiveRelation plaintiff48357 apartment49473) (doneBy have49102 police-officer49060) (isa have49102 Possession)
(topicOfIndividual search49146 apartment49473) (subjectOfMentalSituation search49146 police-officer49060)
(objectActedOn have49102 warrant49216) (isa plaintiff48357 Tenant) (compoundNoun search49146 warrant49216)
(relationExistsInstance defendants CourtCase police-officer49060) (isa apartment49473 ApartmentUnit)
(isa police-officer49060 SocialBeing))
(and (isa investigation49715 CriminalInvestigation) (not (denotes plaintiff48357 investigation49715))
(objectOfInvestigation investigation49715 plaintiff48357) (isa plaintiff48357 SocialBeing)
(relationExistsInstance plaintiffs CourtCase plaintiff48357))
(and (isa police49776 PoliceOfficer-Municipal) (agentCaptured arrest49831 plaintiff48357)
(arrestingOfficer arrest49831 police49776) (isa arrest49831 ArrestingSomeone) (isa plaintiff48357 Tenant))
(and (doneBy hear49967 plaintiff48357) (loudnessLevel noise50024 (HighAmountFn AcousticNoiseLevel))
(topicOfIndividual hear49967 noise50024) (objectFoundInLocation noise50024 yard50132) (isa hear49967 Hearing)
(from-Generic noise50024 plaintiff48357) (isa noise50024 AudibleSound) (isa yard50132 Yard-UnitOfDistance)
(isa plaintiff48357 SocialBeing) (relationExistsInstance plaintiffs CourtCase plaintiff48357))
(and (communicationTarget call50264 |91150310|) (isa call50264 MakingAPhoneCall) (performedBy call50264 plaintiff48357)
(isa |91150310| GovernmentalOrganization) (relationExistsInstance plaintiffs CourtCase plaintiff48357)
(isa plaintiff48357 SocialBeing) (relationExistsInstance nextInSequence List call50264)

```
      (isa |91150310| EmergencyDispatchService))
 (and (performedBy run50398 |91150310|) (endOfPath |91150310| enclose50508)
(startOfPath |91150310| house50440)
      (isa run50398 Running) (from-Generic run50398 house50440) (isa house50440
House-Modern)
      (isa enclose50508 EnclosingSomething) (isa porch50538 Porch)
(relationExistsInstance plaintiffs CourtCase |91150310|)
      (enclosedObject enclose50508 porch50538) (isa |91150310| SocialBeing)
      (relationExistsInstance nextInSequence List run50398))
 (and (isa enter50680 ArrivingAtAPlace) (to-Generic enter50680 porch50538)
(objectMoving enter50680 |91150310|)
      (isa |91150310| SocialBeing) (isa porch50538 Porch)
(relationExistsInstance nextInSequence List enter50680)
      (relationExistsInstance defendants CourtCase |91150310|))
 (and (recipientOfInfo tell50801 |91150310|) (isa tell50801 Informing)
(performedBy tell50801 |91150310|)
      (relationExistsInstance defendants CourtCase |91150310|)
(relationExistsInstance plaintiffs CourtCase |91150310|)
      (isa he50880 PoliceOfficer-Municipal) (isa |91150310| SocialBeing))
 (and (objectFoundInLocation group-of-dog51115 apartment51181) (by-
Underspecified tell50801 group-of-dog51115)
      (relationInstanceExists by-Underspecified tell50801
       (CollectionSubsetFn Dog
        (TheSetOf ?dog51115 (and (isa ?dog51115 Dog) (possessiveRelation
plaintiff51106 group-of-dog51115)))))
      (relationExistsInstance objectFoundInLocation
       (CollectionSubsetFn Dog
        (TheSetOf ?dog51115 (and (isa ?dog51115 Dog) (possessiveRelation
plaintiff51106 group-of-dog51115))))
       apartment51181)
      (isa tell50801 Situation) (isa group-of-dog51115 (SetOfTypeFn Dog)) (isa
apartment51181 ApartmentUnit)
      (possessiveRelation |91150310| apartment51181) (elementOf dog51115 group-
of-dog51115) (isa dog51115 Dog)
      (possessiveRelation plaintiff51106 group-of-dog51115))
 (elementOf dog51115 group-of-dog51115) (and (isa dog51115 Dog)
(possessiveRelation plaintiff51106 group-of-dog51115))
 (and (emitter bark51250 group-of-dog51115) (relationInstanceExists emitter
bark51250 Dog)
      (isa group-of-dog51115 (SetOfTypeFn Dog)) (relationInstanceExists
waveEmitted bark51250 BarkingSound)
      (isa bark51250 MakingAnOralSound) (elementOf dog51228 group-of-dog51115)
(isa dog51228 Dog))
 (elementOf dog51228 group-of-dog51115) (isa dog51228 Dog)
 (and (performedBy run51309 group-of-dog51115) (relationInstanceExists
performedBy run51309 Dog)
      (eventOccursAtLocationType run51309 OutdoorLocation) (isa group-of-
dog51115 (SetOfTypeFn Dog))
      (relationExistsInstance nextInSequence List run51309) (isa run51309
Running) (elementOf dog51298 group-of-dog51115)
      (isa dog51298 Dog))
 (elementOf dog51298 group-of-dog51115) (isa dog51298 Dog)
 (and (cardinality group-of-dog51525 2) (objectActedOn grab51428 group-of-
dog51525)
      (relationInstanceExists objectActedOn grab51428 Dog)
(relationExistsInstance cardinality Dog 2)
```

```
      (rateOfEvent grab51428 (HighToVeryHighAmountFn EventRate))
(relationExistsInstance plaintiffs CourtCase |91150310|)
      (relationExistsInstance nextInSequence List grab51428) (isa |91150310|
SocialBeing)
      (isa grab51428 GrabbingOntoWithHand) (doneBy grab51428 |91150310|) (isa
group-of-dog51525 (SetOfTypeFn Dog))
      (elementOf dog51525 group-of-dog51525) (isa dog51525 Dog))
 (elementOf dog51525 group-of-dog51525) (isa dog51525 Dog)
 (and (objectActedOn hold51672 group-of-dog51727) (relationInstanceExists
objectActedOn hold51672 Dog)
      (relationExistsInstance plaintiffs CourtCase |91150310|) (isa group-of-
dog51727 (SetOfTypeFn Dog))
      (isa |91150310| SocialBeing) (isa hold51672 HoldingAnObject) (isa
chest51806 Chest-BodyPart)
      (to-Generic hold51672 chest51806) (doneBy hold51672 |91150310|)
(possessiveRelation |91150310| chest51806)
      (elementOf dog51727 group-of-dog51727) (isa dog51727 Dog))
 (elementOf dog51727 group-of-dog51727) (isa dog51727 Dog)
 (and (isa dog51969 Dog) (elementOf dog51969 group-of-dog51115) (isa dog51969
Dog) (isa |91150310| SocialBeing)
      (not (and (isa |91150310| SocialBeing) (near group-of-dog51115
|91150310|)
                 (relationExistsInstance defendants CourtCase |91150310|) (near
|91150310| |91150310|)))
      (relationExistsInstance plaintiffs CourtCase |91150310|) (isa group-of-
dog51115 (SetOfTypeFn Dog)))
 (elementOf dog51969 group-of-dog51115) (isa dog51969 Dog)
 (and (isa |91150310| SocialBeing) (near group-of-dog51115 |91150310|)
      (relationExistsInstance defendants CourtCase |91150310|) (near |91150310|
|91150310|))
 (and (threatenedAgent aim-693409 |91150310|) (threateningAgent aim-693409
|91150310|)
      (isa aim-693409 MakingAThreat) (doneBy aim-693409 |91150310|)
      (intendedToLocation aim-693409 |91150310|) (instrument-Generic aim-
693409 weapon-081665)
      (isa aim-693409 AimingSomething)
      (isa hit-689433
       (SubcollectionOfWithRelationToFn
        (SubcollectionOfWithRelationToFn ShootingSomeone objectActedOn
|91150310|) doneBy
        |91150310|))
      (topicOfIndividual anticipation-788347 hit-689433)
      (subjectOfMentalSituation anticipation-788347 |91150310|) (isa
anticipation-788347 Expectation)
      (isa weapon-081665 Weapon) (isa shoot12327 (CausingFn DamageOutcome))
      (isa shoot12327 DestructionEvent) (damages shoot12327 dog51115)
(objectAttacked shoot12327 dog51115)
      (performedBy shoot12327 |91150310|) (isa shoot12327 AttackOnTangible)
      (isa shoot12327 ShootingAndHittingSomething) (objectActedOn hold51672
dog51115)
      (isa hold51672 HoldingAnObject) (doneBy hold51672 |91150310|) (isa
|91150310| Defendant)
      (intendedToLocation shoot12327 dog51115) (isa |91150310| Plaintiff)
      (isa shoot12327 ShootingSomeone) (doneBy shoot12327 |91150310|)
      (relationExistsInstance nextInSequence List shoot12327)
      (isa dog51115 Dog) (occursDuring shoot12327 hold51672))
```

```
(in-microtheory (cleanLanguageOutputMtFn Amos_v_State))
(genlMt (cleanLanguageOutputMtFn Amos_v_State) (CaseLawCorpusMtFn
Amos_v_State))

;;;The plaintiff owned a building.
(doneBy own48367 plaintiff48357) (isa building48405 RealEstate) (isa own48367
OwningSomething)
      (objectActedOn own48367 building48405) (relationExistsInstance
plaintiffs CourtCase plaintiff48357)
      (isa building48405 Building) (isa plaintiff48357 SocialBeing)

;;;The plaintiff lived in the building.
(residence-Role live48468 building48405)
(isa live48468 ResidingSomewhere)
(residents-Role live48468 plaintiff48357)
(isa building48405 Building)

;;;The plaintiff owned dogs.
(doneBy own48368 plaintiff48357)
(isa own48368 OwningSomething)
(objectActedOn own48368 group-of-dog51115)
(isa dog48641 Dog)
(elementOf dog48641 group-of-dog51115)

;;;A tenant lived in an apartment in the building.
(objectFoundInLocation apartment48814 building48405)
(residence-Role live48728 apartment48814)
(isa live48728 ResidingSomewhere)
(objectFoundInLocation tenant48681 apartment48814)
(residents-Role live48728 tenant48681)
(isa apartment48814 ApartmentUnit)
(isa tenant48681 Tenant)

;;;The defendant is a police officer.
(relationExistsInstance defendants CourtCase tenant48681)
(isa tenant48681 SocialBeing)
(isa defendant48982 PoliceOfficer-Municipal)

;;;The defendant had a search warrant for the tenant's apartment.
(isa warrant49216 Examination-Investigation)
(isa warrant49216 Warrant-Writ)
(doneBy have49102 defendant48982)
(isa have49102 Possession)
(topicOfIndividual search49146 apartment48814)
(subjectOfMentalSituation search49146 defendant48982)
(objectActedOn have49102 warrant49216)

;;;The  plaintiff was not under investigation.
(not (and (isa investigation49715 CriminalInvestigation)
            (objectOfInvestigation investigation49715 plaintiff48357)))

;;;The police arrested the tenant.
(isa defendant48982 PoliceOfficer-Municipal)
(agentCaptured arrest49831 tenant48681)
(arrestingOfficer arrest49831 defendant48982)
(isa arrest49831 ArrestingSomeone)
```

```
;;;The plaintiff heard loud noises in the yard.
(doneBy hear49967 plaintiff48357)
(loudnessLevel noise50024 (HighAmountFn AcousticNoiseLevel))
(topicOfIndividual hear49967 noise50024)
(objectFoundInLocation noise50024 yard50132)
(isa hear49967 Hearing)
(from-Generic noise50024 yard50132)
(isa noise50024 AudibleSound)
(isa yard50132 Backyard)

;;;Then the plaintiff called 911.
(relationExistsInstance nextInSequence List call50264)
(communicationTarget call50264 911EmergencyNum)
(isa call50264 MakingAPhoneCall)
(performedBy call50264 plaintiff48357)
(isa 911EmergencyNum GovernmentalOrganization)
(isa 911EmergencyNum EmergencyDispatchService)

;;;Then the plaintiff ran from the house to an enclosed porch.
(relationExistsInstance nextInSequence List run50398)
(performedBy run50398 plaintiff48357)
(endOfPath run50398 porch50538)
(startOfPath run50398 house50440)
(isa run50398 Running)
(from-Generic run50398 building48405)
(isa building48405 House-Modern)
(isa enclose50508 EnclosingSomething)
(isa porch50538 Porch)
(enclosedObject enclose50508 porch50538)

;;;Then the defendant entered the porch.
(relationExistsInstance nextInSequence List enter50680)
(isa enter50680 ArrivingAtAPlace)
(to-Generic enter50680 porch50538)
(objectMoving enter50680 defendant48982)
(isa porch50538 Porch)

;;;The defendant told the plaintiff he was a police officer.
(recipientOfInfo tell50801 plaintiff48357)
(isa tell50801 Informing)
(performedBy tell50801 defendant48982)

;;;The plaintiff's dogs were in her apartment.
(objectFoundInLocation group-of-dog51115 apartment51181)
(isa group-of-dog51115 (SetOfTypeFn Dog))
(isa apartment51181 ApartmentUnit)
(possessiveRelation plaintiff48357 apartment51181)
(possessiveRelation plaintiff48357 group-of-dog51115)

;;;The dogs  were barking.
(emitter bark51250 group-of-dog51115)
(relationInstanceExists waveEmitted bark51250 BarkingSound)
(isa bark51250 MakingAnOralSound)

;;;Then the dogs ran outside.
(performedBy run51309 group-of-dog51115)
(eventOccursAtLocationType run51309 OutdoorLocation)
```

```
(relationExistsInstance nextInSequence List run51309)
(isa run51309 Running)
(elementOf dog51298 group-of-dog51115)
(isa dog51298 Dog)

;;;Then the plaintiff quickly grabbed both dogs.
(relationExistsInstance nextInSequence List grab51428)
(cardinality group-of-dog51115 2)
(objectActedOn grab51428 group-of-dog51115)
(rateOfEvent grab51428 (HighToVeryHighAmountFn EventRate))
(isa grab51428 GrabbingOntoWithHand)
(doneBy grab51428 plaintiff48357)

;;;The plaintiff held the dogs to her chest.
(objectActedOn hold51672 group-of-dog51115)
(isa hold51672 HoldingAnObject)
(isa chest51806 Chest-BodyPart)
(to-Generic hold51672 chest51806)
(doneBy hold51672 plaintiff48357)
(possessiveRelation plaintiff48357 chest51806)

;;;The plaintiff and the dogs were not near the defendant.
(not (and (near group-of-dog51115 defendant48982)
          (near plaintiff48357 defendant48982)))

;;;Then the defendant shot the dog while the plaintiff was holding it.
(startsAfterStartingOf shoot12327 hold51672)
(threatenedAgent aim-693409 plaintiff48357) (threateningAgent aim-693409
defendant48982)
(isa aim-693409 MakingAThreat) (doneBy aim-693409 defendant48982)
(intendedToLocation aim-693409 plaintiff48357) (instrument-Generic aim-693409
weapon-081665)
(isa aim-693409 AimingSomething)
(isa hit-689433
   (SubcollectionOfWithRelationToFn
      (SubcollectionOfWithRelationToFn ShootingSomeone objectActedOn
plaintiff48357) doneBy
      defendant48982))
(topicOfIndividual anticipation-788347 hit-689433)
(subjectOfMentalSituation anticipation-788347 plaintiff48357) (isa
anticipation-788347 Expectation)
(isa weapon-081665 Weapon) (isa shoot12327 (CausingFn DamageOutcome))
(isa shoot12327 DestructionEvent) (damages shoot12327 dog51115) (objectAttacked
shoot12327 dog51115)
(performedBy shoot12327 defendant48982) (isa shoot12327 AttackOnTangible)
(isa shoot12327 ShootingAndHittingSomething) (objectActedOn hold51672 dog51115)
(isa hold51672 HoldingAnObject) (doneBy hold51672 plaintiff48357) (isa
defendant48982 Defendant)
(intendedToLocation shoot12327 dog51115) (isa plaintiff48357 Plaintiff)
(isa shoot12327 ShootingSomeone) (doneBy shoot12327 defendant48982)
(isa dog51115 Dog) (occursDuring shoot12327 hold51672)

(in-microtheory (LegalCaseMtFn Amos_v_State))
(genlMt (LegalCaseMtFn Amos_v_State) (CaseLawCorpusMtFn Amos_v_State))

;;;The plaintiff owned a building.
(isa Amos_v_State CourtCase)
```

```
(plaintiffs Amos_v_State plaintiff48357)
(defendants Amos_v_State defendant48982)
(doneBy own48367 plaintiff48357)
(isa building48405 RealEstate)
(isa own48367 OwningSomething)
(objectActedOn own48367 building48405)
(isa building48405 Building)
(isa plaintiff48357 Plaintiff)

;;;The plaintiff lived in the building.
(residence-Role live48468 building48405)
(isa live48468 ResidingSomewhere)
(residents-Role live48468 plaintiff48357)
(isa building48405 Building)

;;;The plaintiff owned dogs.
(doneBy own48368 plaintiff48357)
(isa own48368 OwningSomething)
(objectActedOn own48368 group-of-dog51115)
(isa dog48641 Dog)
(elementOf dog48641 group-of-dog51115)

;;;A tenant lived in an apartment in the building.
(objectFoundInLocation apartment48814 building48405)
(residence-Role live48728 apartment48814)
(isa live48728 ResidingSomewhere)
(objectFoundInLocation tenant48681 apartment48814)
(residents-Role live48728 tenant48681)
(isa apartment48814 ApartmentUnit)
(isa tenant48681 Tenant)

;;;The defendant is a police officer.
(isa defendant48982 Defendant)
(isa defendant48982 PoliceOfficer-Municipal)

;;;The defendant had a search warrant for the tenant's apartment.
(isa warrant49216 Warrant-Writ)
(doneBy have49102 defendant48982)
(isa have49102 Possession)
(topicOfIndividual search49146 apartment48814)
(subjectOfMentalSituation search49146 defendant48982)
(objectActedOn have49102 warrant49216)

;;;The  plaintiff was not under investigation.
(not (and (isa investigation49715 CriminalInvestigation)
            (objectOfInvestigation investigation49715 plaintiff48357)))

;;;The police arrested the tenant.
(isa defendant48982 PoliceOfficer-Municipal)
(agentCaptured arrest49831 tenant48681)
(arrestingOfficer arrest49831 defendant48982)
(isa arrest49831 ArrestingSomeone)

;;;The plaintiff heard loud noises in the yard.
(doneBy hear49967 plaintiff48357)
(loudnessLevel noise50024 (HighAmountFn AcousticNoiseLevel))
(topicOfIndividual hear49967 noise50024)
```

```
(objectFoundInLocation noise50024 yard50132)
(isa hear49967 Hearing)
(from-Generic noise50024 yard50132)
(isa noise50024 AudibleSound)
(isa yard50132 Backyard)

;;;Then the plaintiff called 911.
(startsAfterStartingOf call50264 hear49967)
(communicationTarget call50264 911EmergencyNum)
(isa call50264 MakingAPhoneCall)
(performedBy call50264 plaintiff48357)
(isa 911EmergencyNum GovernmentalOrganization)
(isa 911EmergencyNum EmergencyDispatchService)

;;;Then the plaintiff ran from the house to an enclosed porch.
(startsAfterStartingOf run50398 call50264)
(performedBy run50398 plaintiff48357)
(endOfPath run50398 porch50538)
(startOfPath run50398 building48405)
(isa run50398 Running)
(from-Generic run50398 building48405)
(isa building48405 House-Modern)
(isa enclose50508 EnclosingSomething)
(isa porch50538 Porch)
(enclosedObject enclose50508 porch50538)

;;;Then the defendant entered the porch.
(startsAfterStartingOf enter50680 run50398)
(isa enter50680 ArrivingAtAPlace)
(to-Generic enter50680 porch50538)
(objectMoving enter50680 defendant48982)
(isa porch50538 Porch)

;;;The defendant told the plaintiff he was a police officer.
(recipientOfInfo tell50801 plaintiff48357)
(isa tell50801 Informing)
(performedBy tell50801 defendant48982)
(topicOfInfoTransfer tell50801 (isa defendant48982 PoliceOfficer-Municipal))

;;;The plaintiff's dogs were in her apartment.
(objectFoundInLocation group-of-dog51115 apartment51181)
(isa group-of-dog51115 (SetOfTypeFn Dog))
(isa dog48641 Dog) (elementOf dog48641 group-of-dog51115)
(isa apartment51181 ApartmentUnit)
(possessiveRelation plaintiff48357 apartment51181)
(possessiveRelation plaintiff48357 group-of-dog51115)

;;;The dogs  were barking.
(emitter bark51250 group-of-dog51115)
(relationInstanceExists waveEmitted bark51250 BarkingSound)
(isa bark51250 MakingAnOralSound)

;;;Then the dogs ran outside.
(startsAfterStartingOf run51309 bark51250)
(performedBy run51309 group-of-dog51115)
(eventOccursAtLocationType run51309 OutdoorLocation)
(toLocation run51309 yard50132)
```

```
(isa run51309 Running)

;;;Then the plaintiff quickly grabbed both dogs.
(relationExistsInstance nextInSequence List grab51428)
(cardinality group-of-dog51115 2)
(objectActedOn grab51428 group-of-dog51115)
(rateOfEvent grab51428 (HighToVeryHighAmountFn EventRate))
(isa grab51428 GrabbingOntoWithHand)
(doneBy grab51428 plaintiff48357)

;;;The plaintiff held the dogs to her chest.
(objectActedOn hold51672 group-of-dog51115)
(isa hold51672 HoldingAnObject)
(isa chest51806 Chest-BodyPart)
(to-Generic hold51672 chest51806)
(doneBy hold51672 plaintiff48357)
(possessiveRelation plaintiff48357 chest51806)

;;;The plaintiff and the dogs were not near the defendant.
(not (and (near group-of-dog51115 defendant48982)
          (near plaintiff48357 defendant48982)))

;;;Then the defendant shot the dog while the plaintiff was holding it.
(startsAfterStartingOf shoot12327 hold51672)
(threatenedAgent aim-693409 plaintiff48357) (threateningAgent aim-693409
defendant48982)
(isa aim-693409 MakingAThreat) (doneBy aim-693409 defendant48982)
(intendedToLocation aim-693409 plaintiff48357) (instrument-Generic aim-693409
weapon-081665)
(isa aim-693409 AimingSomething)
(isa hit-689433
   (SubcollectionOfWithRelationToFn
      (SubcollectionOfWithRelationToFn ShootingSomeone objectActedOn
plaintiff48357) doneBy
      defendant48982))
(topicOfIndividual anticipation-788347 hit-689433)
(subjectOfMentalSituation anticipation-788347 plaintiff48357) (isa
anticipation-788347 Expectation)
(isa weapon-081665 Weapon) (isa shoot12327 (CausingFn DamageOutcome))
(isa shoot12327 DestructionEvent) (damages shoot12327 dog48641) (objectAttacked
shoot12327 dog48641)
(performedBy shoot12327 defendant48982) (isa shoot12327 AttackOnTangible)
(isa shoot12327 ShootingAndHittingSomething) (objectActedOn hold51672 dog48641)
(isa hold51672 HoldingAnObject) (doneBy hold51672 plaintiff48357)
(intendedToLocation shoot12327 dog48641) (occursDuring shoot12327 hold51672)
(isa shoot12327 ShootingSomeone) (doneBy shoot12327 defendant48982)

(in-microtheory (LegalCaseConclusionMtFn Amos_v_State))
(genlMt (LegalCaseConclusionMtFn Amos_v_State) (LegalCaseMtFn Amos_v_State))

(battersPartyByDoing defendant48982 plaintiff48357 shoot52213)
(assaultsPartyByDoing defendant48982 plaintiff48357 shoot52213)

(in-microtheory (LegalCaseConclusion-NegatedMtFn Amos_v_State))
(not (battersPartyByDoing defendant48982 plaintiff48357 shoot52213))
(not (assaultsPartyByDoing defendant48982 plaintiff48357 shoot52213))
```

```
(in-microtheory (LegalCaseConclusion-ReversedMtFn Amos_v_State))
(battersPartyByDoing plaintiff48357 defendant48982 shoot52213)
(assaultsPartyByDoing plaintiff48357 defendant48982 shoot52213)

(in-microtheory (LegalCaseConclusion-NegatedReversedMtFn Amos_v_State))
(not (battersPartyByDoing plaintiff48357 defendant48982 shoot52213))
(not (assaultsPartyByDoing plaintiff48357 defendant48982 shoot52213))
```

(in-microtheory CaseLawCorpusMt)

(isa Ariola_v_Nigro CaseLawCorpusCase)

(in-microtheory (CaseLawCorpusMtFn Ariola_v_Nigro))

(caseType Ariola_v_Nigro Trespass)

(caseValence Ariola_v_Nigro Trespass Positive)

(caseName Ariola_v_Nigro "Ariola v. Nigro")

(caseCourt Ariola_v_Nigro "Supreme Court of Illinois")

(caseYear Ariola_v_Nigro "1959")

(caseReporter Ariola_v_Nigro "16 Ill. 2d 46")

(caseOriginalText Ariola_v_Nigro "Since June 9, 1925, plaintiffs Saverio Ariola and Susanna Ariola were owners in joint tenancy of the property at 818 N. Twenty-third Avenue, Melrose Park, which was improved with a two-story house occupied by them and their children, the other plaintiffs herein, as a family home. The Ariola property fronts on Twenty-third Avenue and extends west for a distance of approximately 50 feet along Iowa Street on the north, and for 51 feet along Lake Street on the south. Immediately to the west of the Ariola property is the Nigro property, held in joint tenancy by defendants. This property, known as 2305 Lake Street, fronts on Lake Street, and extends along that street for some 77 feet 3 1/2 inches, and along Iowa Street on the north for some 75 feet 1 3/4 inches. The property was improved with a one-story building, until October, 1948, when defendants commenced construction of an addition. According to the testimony of defendants' mason contractor, he excavated right up to plaintiffs' building foundation, put in forms only on one side and poured concrete flush against plaintiffs' foundation, so that the east part of defendants' foundation is flush with the west part of plaintiffs' building. Special surveys made at plaintiffs' request by registered surveyors before defendants' brick work was superimposed, indicated that defendants' foundation encroached upon plaintiffs' property to the extent of 1 inch at the northeast corner of defendants' foundation and some 2 3/8 inches at the southeast corner of the foundation. Plaintiffs thereupon notified defendants of the encroachment and requested them to discontinue construction. Notice was also given by plaintiffs' attorney to defendant Nigro and to his attorney. Nevertheless, defendants proceeded with the construction of the two-story brick addition, which according to plaintiffs' surveys also encroached to the same extent above ground level. Defendants, however, deny any such encroachment, and their mason contractor testified that the village markers were followed, and that plaintiffs' foundation was irregular and encroached on defendants' property. It appears further from defendant Daniel Nigro's testimony as an adverse witness that before construction began he had notified plaintiffs that defendants would occupy all of their property, and that it would be necessary for plaintiffs to remove their projecting gutters and downspouts along the west wall. The evidence is controverted as to the length of time plaintiffs' gutters and downspouts projected from plaintiff's west wall. On plaintiffs' behalf, their original builder and other witnesses testified that the gutters and downspouts were originally installed along the west wall in 1925, when the house was built, and have remained there since that time. On defendants' behalf, his former partner testified that the gutters were not there in 1945, that there had been gutters on the south wall, as suggested by a photograph showing some discoloration on the south wall, possibly from rain gutters; and the brother of defendants' attorney stated that as plaintiffs' agent he had secured permission from defendants to install the gutters about a year and onehalf before they were torn down. However, a plat of survey, introduced by defendants, and dated June 7, 1941, contained the notation, 'downspouts and gutters project 6' west of line.' In any event, plaintiffs refused to remove the gutters, and, according to defendants' admission and the testimony of their builder, defendants ordered that they be torn down and that the construction of the building proceed as planned. The evidence further shows that upon plaintiffs' refusal to pay half the cost of installing a 'saddle' type drainage installation, defendants, unbeknown to plaintiffs, installed a tar paper flashing between the east wall of the new building and the west wall of plaintiffs' home, to prevent drainage of rain water and melting snow between the two buildings. This flashing, according to the testimony of both plaintiffs' and defendants' witnesses, was not very effective and caused the accumulation of water along the west wall of the Ariola building, with resulting seepage and rotting of the plaster flashing on the inside of plaintiffs' wall and deterioration of the mortar joints between the bricks. Plaintiffs offered evidence that such seepage did not occur prior to 1948. Upon discovery of the tar paper flashing in 1951, some two years after it was installed, plaintiff Ariola's son kicked a hole in it in order to permit the water to drain off. His testimony as to the accumulation of several inches of water at the west end of the roof was corroborated by the commercial photographer who saw the premises in 1951.")

(caseConclusion Ariola_v_Nigro "Plaintiffs have an easement through adverse possession for the maintenance of their gutters")

(caseConclusion Ariola_v_Nigro "Defendants encroached on plaintiffs' property")

(caseSimpleText Ariola_v_Nigro "The plaintiffs own a house. The plaintiffs' property is adjacent to the defendants' property. The defendants began to build an addition. The plaintiffs' gutters were over the defendants' property. The gutters had existed for twenty years. The defendants told the plaintiffs to remove the gutters. The defendants built a new foundation. The defendants' foundation was on on the plaintiffs' property. The plaintiffs notified the defendants of this. The plaintiffs asked the defendants to stop building the addition. The defendants destroyed the plaintiffs' gutters. The defendants built the addition. The defendants installed a drainage system. The drainage system did not function. The drainage system damaged the plaintiffs' building.")

```
(in-microtheory (rawLanguageOutputMtFn Ariola_v_Nigro))
(genlMt (rawLanguageOutputMtFn Ariola_v_Nigro) (CaseLawCorpusMtFn
Ariola_v_Nigro))

(and (doneBy own74910 group-of-plaintiff74903)
      (relationInstanceExists doneBy own74910
       (CollectionSubsetFn SocialBeing
        (TheSetOf ?plaintiff74903
         (and (relationExistsInstance plaintiffs CourtCase ?plaintiff74903)
(isa ?plaintiff74903 SocialBeing)))))
      (isa group-of-plaintiff74903 (SetOfTypeFn SocialBeing)) (objectActedOn
own74910 house74952)
      (isa own74910 OwningSomething) (isa house74952 House-Modern) (elementOf
plaintiff74903 group-of-plaintiff74903)
      (relationExistsInstance plaintiffs CourtCase plaintiff74903) (isa
plaintiff74903 SocialBeing))
 (elementOf plaintiff74903 group-of-plaintiff74903)
```

```
 (and (relationExistsInstance plaintiffs CourtCase plaintiff74903) (isa
plaintiff74903 SocialBeing))
 (and (isa property75133 Property) (owns group-of-defendant75101 property75133)
(isa property75047 Property)
      (owns group-of-plaintiff75036 property75047) (endOfPath property75047
property75133)
      (isa group-of-plaintiff75036 (SetOfTypeFn SocialBeing))
(possessiveRelation group-of-defendant75101 property75133)
      (isa group-of-defendant75101 (SetOfTypeFn SocialBeing))
(possessiveRelation group-of-plaintiff75036 property75047))
 (and (doneBy build75255 group-of-plaintiff74903) (doneBy begin75234 group-of-
plaintiff74903)
      (relationInstanceExists doneBy build75255
       (CollectionSubsetFn SocialBeing
        (TheSetOf ?defendant75226
         (and (relationExistsInstance defendants CourtCase ?defendant75226)
(isa ?defendant75226 SocialBeing)))))
      (relationInstanceExists doneBy begin75234
       (CollectionSubsetFn SocialBeing
        (TheSetOf ?defendant75226
         (and (relationExistsInstance defendants CourtCase ?defendant75226)
(isa ?defendant75226 SocialBeing)))))
      (isa group-of-plaintiff74903 (SetOfTypeFn SocialBeing)) (isa
addition75315 NewArtifact)
      (isa build75255 MakingSomething) (situationBeginning begin75234
build75255) (outputsCreated build75255 addition75315)
      (isa begin75234 BeginningAnActivity) (isa addition75315
RoomInAConstruction)
      (elementOf defendant75226 group-of-plaintiff74903)
(relationExistsInstance defendants CourtCase defendant75226)
      (isa defendant75226 SocialBeing))
 (elementOf defendant75226 group-of-plaintiff74903)
 (and (relationExistsInstance defendants CourtCase defendant75226) (isa
defendant75226 SocialBeing))
 (and (objectActedOn own74910 group-of-gutter78692) (relationInstanceExists
objectActedOn own74910 GutterOnARoof)
      (measure own74910 (YearsDuration 20)) (isa group-of-gutter78692
(SetOfTypeFn GutterOnARoof))
      (fe_viewpoint own74910 year75672) (isa own74910 Subsist-
SomethingExisting) (elementOf gutter75559 group-of-gutter78692)
      (isa gutter75559 GutterOnARoof))
 (elementOf gutter75559 group-of-gutter78692) (isa gutter75559 GutterOnARoof)
 (and (relationInstanceExists objectRemoved remove75816 GutterOnARoof)
      (isa group-of-plaintiff75785 (SetOfTypeFn SocialBeing)) (isa group-of-
defendant75723 (SetOfTypeFn SocialBeing))
      (isa group-of-gutter75946 (SetOfTypeFn GutterOnARoof)) (isa tell75728
Requesting-CommunicationAct)
      (recipientOfInfo tell75728 group-of-plaintiff75785) (isa remove75816
RemovingSomething)
      (performedBy tell75728 group-of-defendant75723) (situationTopic
tell75728 remove75816)
      (doneBy remove75816 group-of-plaintiff75785) (elementOf gutter75946
group-of-gutter75946)
      (isa gutter75946 GutterOnARoof)
      (relationInstanceExists performedBy tell75728
       (CollectionSubsetFn SocialBeing
        (TheSetOf ?defendant75723
```

```
        (and (relationExistsInstance defendants CourtCase ?defendant75723)
(isa ?defendant75723 SocialBeing)))))
      (isa group-of-plaintiff75785 (SetOfTypeFn SocialBeing)) (isa group-of-
defendant75723 (SetOfTypeFn SocialBeing))
      (isa group-of-gutter75946 (SetOfTypeFn GutterOnARoof)) (isa tell75728
Requesting-CommunicationAct)
      (objectRemoved remove75816 group-of-gutter75946) (recipientOfInfo
tell75728 group-of-plaintiff75785)
      (isa remove75816 RemovingSomething) (situationTopic tell75728
remove75816)
      (doneBy remove75816 group-of-plaintiff75785) (elementOf defendant75723
group-of-defendant75723)
      (relationExistsInstance defendants CourtCase defendant75723) (isa
defendant75723 SocialBeing)
      (relationInstanceExists doneBy remove75816
       (CollectionSubsetFn SocialBeing
        (TheSetOf ?plaintiff75785
         (and (relationExistsInstance plaintiffs CourtCase ?plaintiff75785)
(isa ?plaintiff75785 SocialBeing)))))
      (relationInstanceExists recipientOfInfo tell75728
       (CollectionSubsetFn SocialBeing
        (TheSetOf ?plaintiff75785
         (and (relationExistsInstance plaintiffs CourtCase ?plaintiff75785)
(isa ?plaintiff75785 SocialBeing)))))
      (isa group-of-plaintiff75785 (SetOfTypeFn SocialBeing)) (isa group-of-
defendant75723 (SetOfTypeFn SocialBeing))
      (isa group-of-gutter75946 (SetOfTypeFn GutterOnARoof)) (isa tell75728
Requesting-CommunicationAct)
      (objectRemoved remove75816 group-of-gutter75946) (isa remove75816
RemovingSomething)
      (performedBy tell75728 group-of-defendant75723) (situationTopic
tell75728 remove75816)
      (elementOf plaintiff75785 group-of-plaintiff75785)
(relationExistsInstance plaintiffs CourtCase plaintiff75785)
      (isa plaintiff75785 SocialBeing))
 (elementOf defendant75723 group-of-defendant75723)
 (and (relationExistsInstance defendants CourtCase defendant75723) (isa
defendant75723 SocialBeing))
 (elementOf plaintiff75785 group-of-plaintiff75785)
 (and (relationExistsInstance plaintiffs CourtCase plaintiff75785) (isa
plaintiff75785 SocialBeing))
 (elementOf gutter75946 group-of-gutter75946) (isa gutter75946 GutterOnARoof)
 (and (isa plaintiff6974 SocialBeing) (relationExistsInstance plaintiffs
CourtCase plaintiff6974)
      (elementOf plaintiff6974 group-of-plaintiff6974)
      (relationExistsInstance owns
       (CollectionSubsetFn SocialBeing
        (TheSetOf ?plaintiff6974
         (and (relationExistsInstance plaintiffs CourtCase ?plaintiff6974)
              (isa ?plaintiff6974 SocialBeing))))
       property7022)
      (relationExistsInstance possessiveRelation
       (CollectionSubsetFn SocialBeing
        (TheSetOf ?plaintiff6974
         (and (relationExistsInstance plaintiffs CourtCase ?plaintiff6974)
              (isa ?plaintiff6974 SocialBeing))))
       property7022)
```

```
      (isa defendant6872 SocialBeing) (relationExistsInstance defendants
CourtCase defendant6872)
      (elementOf defendant6872 group-of-defendant6872)
      (relationInstanceExists doneBy build75255
       (CollectionSubsetFn SocialBeing
        (TheSetOf ?defendant6872
         (and (relationExistsInstance defendants CourtCase ?defendant6872)
              (isa ?defendant6872 SocialBeing)))))
      (isa group-of-defendant6872 (SetOfTypeFn SocialBeing))
      (isa group-of-plaintiff6974 (SetOfTypeFn SocialBeing)) (isa build75255
BuildingSomething-Supervising)
      (doneBy build75255 group-of-defendant6872) (outputsCreated build75255
foundation6927)
      (eventOccursAt build75255 property7022)
      (causes-PropProp
       (and (isa build75255 BuildingSomething-Supervising) (doneBy build75255
group-of-defendant6872))
       (objectFoundInLocation foundation6927 property7022))
      (owns group-of-plaintiff6974 property7022) (isa property7022 Property)
      (isa foundation6927 FoundationOfABuilding) (possessiveRelation group-of-
plaintiff6974 property7022))
 (elementOf defendant76189 group-of-defendant76189)
 (and (isa defendant76189 SocialBeing) (relationExistsInstance defendants
CourtCase defendant76189))
 (elementOf plaintiff76268 group-of-plaintiff76268)
 (and (relationExistsInstance plaintiffs CourtCase plaintiff76268) (isa
plaintiff76268 SocialBeing))
 (and (infoTransferred notify76364
       (and (relationExistsInstance possessiveRelation
             (CollectionSubsetFn SocialBeing
              (TheSetOf ?defendant76189
               (and (isa ?defendant76189 SocialBeing) (relationExistsInstance
defendants CourtCase ?defendant76189))))
             foundation76198)
            (stateOfDevice foundation76198 Device-On) (isa group-of-
defendant76189 (SetOfTypeFn SocialBeing))
            (possessiveRelation group-of-plaintiff76268 property76290)
            (isa group-of-plaintiff76268 (SetOfTypeFn SocialBeing)) (on-
UnderspecifiedSurface foundation76198 property76290)
            (isa foundation76198 FoundationOfABuilding) (isa property76290
Property)
            (owns group-of-plaintiff76268 property76290) (elementOf
defendant76189 group-of-defendant76189)
            (isa defendant76189 SocialBeing) (relationExistsInstance
defendants CourtCase defendant76189)
            (relationExistsInstance owns
             (CollectionSubsetFn SocialBeing
              (TheSetOf ?plaintiff76268
               (and (relationExistsInstance plaintiffs CourtCase
?plaintiff76268) (isa ?plaintiff76268 SocialBeing))))
             property76290)
            (relationExistsInstance possessiveRelation
             (CollectionSubsetFn SocialBeing
              (TheSetOf ?plaintiff76268
               (and (relationExistsInstance plaintiffs CourtCase
?plaintiff76268) (isa ?plaintiff76268 SocialBeing))))
             property76290)
```

```
            (stateOfDevice   foundation76198   Device-On)   (possessiveRelation
group-of-defendant76189 foundation76198)
            (isa group-of-defendant76189 (SetOfTypeFn SocialBeing)) (isa group-
of-plaintiff76268 (SetOfTypeFn SocialBeing))
            (on-UnderspecifiedSurface   foundation76198   property76290)   (isa
foundation76198 FoundationOfABuilding)
            (isa property76290 Property) (elementOf plaintiff76268 group-of-
plaintiff76268)
            (relationExistsInstance plaintiffs CourtCase plaintiff76268) (isa
plaintiff76268 SocialBeing)))
      (recipientOfInfo notify76364 group-of-defendant76420)
      (isa
       (and (relationExistsInstance possessiveRelation
             (CollectionSubsetFn SocialBeing
              (TheSetOf ?defendant76189
               (and (isa ?defendant76189 SocialBeing) (relationExistsInstance
defendants CourtCase ?defendant76189))))
             foundation76198)
            (stateOfDevice   foundation76198   Device-On)   (isa   group-of-
defendant76189 (SetOfTypeFn SocialBeing))
            (possessiveRelation group-of-plaintiff76268 property76290)
            (isa  group-of-plaintiff76268  (SetOfTypeFn  SocialBeing))  (on-
UnderspecifiedSurface foundation76198 property76290)
            (isa  foundation76198  FoundationOfABuilding)  (isa  property76290
Property)
            (owns    group-of-plaintiff76268    property76290)    (elementOf
defendant76189 group-of-defendant76189)
            (isa    defendant76189    SocialBeing)    (relationExistsInstance
defendants CourtCase defendant76189)
            (relationExistsInstance owns
             (CollectionSubsetFn SocialBeing
              (TheSetOf ?plaintiff76268
               (and     (relationExistsInstance     plaintiffs     CourtCase
?plaintiff76268) (isa ?plaintiff76268 SocialBeing))))
             property76290)
            (relationExistsInstance possessiveRelation
             (CollectionSubsetFn SocialBeing
              (TheSetOf ?plaintiff76268
               (and     (relationExistsInstance     plaintiffs     CourtCase
?plaintiff76268) (isa ?plaintiff76268 SocialBeing))))
             property76290)
            (stateOfDevice   foundation76198   Device-On)   (possessiveRelation
group-of-defendant76189 foundation76198)
            (isa group-of-defendant76189 (SetOfTypeFn SocialBeing)) (isa group-
of-plaintiff76268 (SetOfTypeFn SocialBeing))
            (on-UnderspecifiedSurface   foundation76198   property76290)   (isa
foundation76198 FoundationOfABuilding)
            (isa property76290 Property) (elementOf plaintiff76268 group-of-
plaintiff76268)
            (relationExistsInstance plaintiffs CourtCase plaintiff76268) (isa
plaintiff76268 SocialBeing))
       Backreference-ClassA-NLAttr)
      (performedBy   notify76364   group-of-plaintiff76355)   (isa   group-of-
defendant76420 Set-Mathematical)
      (isa   group-of-plaintiff76355   Set-Mathematical)   (isa   notify76364
Informing))
 (and (relationInstanceExists performedBy ask76592
```

```
       (CollectionSubsetFn SocialBeing
        (TheSetOf ?plaintiff76583
         (and (isa ?plaintiff76583 SocialBeing) (relationExistsInstance
plaintiffs CourtCase ?plaintiff76583)))))
      (isa group-of-plaintiff76583 (SetOfTypeFn SocialBeing)) (purposeOf-
Generic build76780 stop76682)
      (isa ask76592 Requesting-CommunicationAct) (isa addition77219
NewArtifact)
      (recipientOfInfo ask76592 group-of-defendant76648)
(temporalThingTerminated stop76682 build76780)
      (doneBy build76780 group-of-defendant76648) (doneBy stop76682 group-of-
defendant76648)
      (situationTopic ask76592 stop76682) (outputsCreated build76780
addition77219) (isa stop76682 DiscontinuingAnActivity)
      (isa group-of-defendant76648 (SetOfTypeFn SocialBeing)) (isa build76780
BuildingSomething-Supervising)
      (isa addition77219 RoomInAConstruction) (elementOf plaintiff76583 group-
of-plaintiff76583)
      (isa plaintiff76583 SocialBeing) (relationExistsInstance plaintiffs
CourtCase plaintiff76583)
      (relationInstanceExists doneBy stop76682
       (CollectionSubsetFn SocialBeing
        (TheSetOf ?defendant76648
         (and (isa ?defendant76648 SocialBeing) (relationExistsInstance
defendants CourtCase ?defendant76648)))))
      (relationInstanceExists doneBy build76780
       (CollectionSubsetFn SocialBeing
        (TheSetOf ?defendant76648
         (and (isa ?defendant76648 SocialBeing) (relationExistsInstance
defendants CourtCase ?defendant76648)))))
      (relationInstanceExists recipientOfInfo ask76592
       (CollectionSubsetFn SocialBeing
        (TheSetOf ?defendant76648
         (and (isa ?defendant76648 SocialBeing) (relationExistsInstance
defendants CourtCase ?defendant76648)))))
      (isa group-of-plaintiff76583 (SetOfTypeFn SocialBeing)) (purposeOf-
Generic build76780 stop76682)
      (isa ask76592 Requesting-CommunicationAct) (isa addition77219
NewArtifact)
      (temporalThingTerminated stop76682 build76780) (performedBy ask76592
group-of-plaintiff76583)
      (situationTopic ask76592 stop76682) (outputsCreated build76780
addition77219) (isa stop76682 DiscontinuingAnActivity)
      (isa group-of-defendant76648 (SetOfTypeFn SocialBeing)) (isa build76780
BuildingSomething-Supervising)
      (isa addition77219 RoomInAConstruction) (elementOf defendant76648 group-
of-defendant76648)
      (isa defendant76648 SocialBeing) (relationExistsInstance defendants
CourtCase defendant76648))
 (elementOf plaintiff76583 group-of-plaintiff76583)
 (and (isa plaintiff76583 SocialBeing) (relationExistsInstance plaintiffs
CourtCase plaintiff76583))
 (elementOf defendant76648 group-of-defendant76648)
 (and (isa defendant76648 SocialBeing) (relationExistsInstance defendants
CourtCase defendant76648))
 (and (doneBy build77993 group-of-defendant75723)
      (relationInstanceExists doneBy build77993
```

```
       (CollectionSubsetFn SocialBeing
        (TheSetOf ?defendant77986
         (and (relationExistsInstance defendants CourtCase ?defendant77986)
(isa ?defendant77986 SocialBeing)))))
      (isa group-of-defendant75723 (SetOfTypeFn SocialBeing)) (isa
addition78048 NewArtifact)
      (isa addition78048 RoomInAConstruction) (isa build77993
BuildingSomething-Supervising)
      (outputsCreated build77993 addition78048) (elementOf defendant77986
group-of-defendant75723)
      (relationExistsInstance defendants CourtCase defendant77986) (isa
defendant77986 SocialBeing))
 (elementOf defendant77986 group-of-defendant75723)
 (and (relationExistsInstance defendants CourtCase defendant77986) (isa
defendant77986 SocialBeing))
 (and (doneBy install78111 group-of-defendant75723)
      (relationInstanceExists doneBy install78111
       (CollectionSubsetFn SocialBeing
        (TheSetOf ?defendant78102
         (and (isa ?defendant78102 SocialBeing) (relationExistsInstance
defendants CourtCase ?defendant78102)))))
      (isa group-of-defendant75723 (SetOfTypeFn SocialBeing)) (isa drainage-
system78178 DrainageSystem)
      (objectActedOn install78111 drainage-system78178) (isa install78111
InstallingAnObject)
      (elementOf defendant78102 group-of-defendant75723) (isa defendant78102
SocialBeing)
      (relationExistsInstance defendants CourtCase defendant78102))
 (elementOf defendant78102 group-of-defendant75723)
 (and (isa defendant78102 SocialBeing) (relationExistsInstance defendants
CourtCase defendant78102))
 (and (possessiveRelation group-of-defendant76066 building78624)
      (relationExistsInstance possessiveRelation
       (CollectionSubsetFn SocialBeing
        (TheSetOf ?plaintiff78598
         (and (relationExistsInstance plaintiffs CourtCase ?plaintiff78598)
(isa ?plaintiff78598 SocialBeing))))
       building78624)
      (isa group-of-defendant76066 (SetOfTypeFn SocialBeing)) (isa drainage-
system78178 DrainageSystem)
      (isa building78624 Building) (isa damage78543 IncurringPhysicalDamage)
(objectActedOn damage78543 building78624)
      (doneBy damage78543 drainage-system78178) (elementOf plaintiff78598
group-of-defendant76066)
      (relationExistsInstance plaintiffs CourtCase plaintiff78598) (isa
plaintiff78598 SocialBeing))
 (elementOf plaintiff78598 group-of-defendant76066)
 (and (relationExistsInstance plaintiffs CourtCase plaintiff78598) (isa
plaintiff78598 SocialBeing))
 (and (relationInstanceExists by-Underspecified own74910
       (CollectionSubsetFn GutterOnARoof
        (TheSetOf ?gutter78692
         (and (isa ?gutter78692 GutterOnARoof) (possessiveRelation group-of-
plaintiff78683 group-of-gutter78692)))))
      (relationExistsInstance over-UnderspecifiedLocation
       (CollectionSubsetFn GutterOnARoof
        (TheSetOf ?gutter78692
```

```
        (and (isa ?gutter78692 GutterOnARoof) (possessiveRelation group-of-
plaintiff78683 group-of-gutter78692))))
      property78763)
     (isa own74910 Situation) (isa group-of-plaintiff74903 (SetOfTypeFn
SocialBeing))
     (isa group-of-gutter78692 (SetOfTypeFn GutterOnARoof))
(possessiveRelation group-of-plaintiff74903 property78763)
     (owns group-of-plaintiff74903 property78763) (isa property78763 Property)
(elementOf gutter78692 group-of-gutter78692)
     (isa gutter78692 GutterOnARoof) (possessiveRelation group-of-
plaintiff78683 group-of-gutter78692)
     (relationExistsInstance owns
      (CollectionSubsetFn SocialBeing
       (TheSetOf ?defendant78748
        (and (relationExistsInstance defendants CourtCase ?defendant78748)
(isa ?defendant78748 SocialBeing))))
      property78763)
     (relationExistsInstance possessiveRelation
      (CollectionSubsetFn SocialBeing
       (TheSetOf ?defendant78748
        (and (relationExistsInstance defendants CourtCase ?defendant78748)
(isa ?defendant78748 SocialBeing))))
      property78763)
     (over-UnderspecifiedLocation group-of-gutter78692 property78763) (by-
Underspecified own74910 group-of-gutter78692)
     (isa own74910 Situation) (isa group-of-plaintiff74903 (SetOfTypeFn
SocialBeing))
     (isa group-of-gutter78692 (SetOfTypeFn GutterOnARoof)) (isa property78763
Property)
     (elementOf defendant78748 group-of-plaintiff74903)
(relationExistsInstance defendants CourtCase defendant78748)
     (isa defendant78748 SocialBeing))
 (elementOf gutter78692 group-of-gutter78692)
 (and (isa gutter78692 GutterOnARoof) (possessiveRelation group-of-
plaintiff78683 group-of-gutter78692))
 (elementOf defendant78748 group-of-plaintiff74903)
 (and (relationExistsInstance defendants CourtCase defendant78748) (isa
defendant78748 SocialBeing))
 (and (relationInstanceExists objectActedOn destroy78828
      (CollectionSubsetFn GutterOnARoof
       (TheSetOf ?gutter78909
        (and (isa ?gutter78909 GutterOnARoof) (possessiveRelation group-of-
plaintiff78883 group-of-gutter78909)))))
     (isa group-of-gutter78909 (SetOfTypeFn GutterOnARoof)) (isa group-of-
defendant75723 (SetOfTypeFn SocialBeing))
     (doneBy destroy78828 group-of-defendant75723) (isa destroy78828
DestructionEvent)
     (elementOf gutter78909 group-of-gutter78909) (isa gutter78909
GutterOnARoof)
     (possessiveRelation group-of-plaintiff78883 group-of-gutter78909)
     (relationInstanceExists doneBy destroy78828
      (CollectionSubsetFn SocialBeing
       (TheSetOf ?defendant78814
        (and (relationExistsInstance defendants CourtCase ?defendant78814)
(isa ?defendant78814 SocialBeing)))))
     (isa group-of-gutter78909 (SetOfTypeFn GutterOnARoof)) (isa group-of-
defendant75723 (SetOfTypeFn SocialBeing))
```

```
      (objectActedOn  destroy78828  group-of-gutter78909)  (isa  destroy78828
DestructionEvent)
      (elementOf          defendant78814          group-of-defendant75723)
(relationExistsInstance defendants CourtCase defendant78814)
      (isa defendant78814 SocialBeing))
 (elementOf defendant78814 group-of-defendant75723)
 (and  (relationExistsInstance  defendants  CourtCase  defendant78814)  (isa
defendant78814 SocialBeing))
 (elementOf gutter78909 group-of-gutter78909)
 (and   (isa   gutter78909   GutterOnARoof)   (possessiveRelation   group-of-
plaintiff78883 group-of-gutter78909))
 (not  (and  (hasEvaluativeQuantity  drainage-system79564  (PositiveAmountFn
Goodness-Functional))
                  (isa drainage-system79564 FunctionalObject)))


;;;;;;;;;;;;;;;;;;;
(in-microtheory (cleanLanguageOutputMtFn Ariola_v_Nigro))
(genlMt    (cleanLanguageOutputMtFn    Ariola_v_Nigro)    (CaseLawCorpusMtFn
Ariola_v_Nigro))

;;;The plaintiffs own a house.
(doneBy own74910 group-of-plaintiff74903)
      (isa group-of-plaintiff74903 (SetOfTypeFn SocialBeing)) (objectActedOn
own74910 house74952)
      (isa own74910 OwningSomething) (isa house74952 House-Modern) (elementOf
plaintiff74903 group-of-plaintiff74903)
      (relationExistsInstance  plaintiffs  CourtCase  plaintiff74903)  (isa
plaintiff74903 SocialBeing)

;;;The plaintiffs' property is adjacent to the defendants' property.
;;; add the plaintiffs own!
(isa property75133 Property) (owns group-of-defendant75101 property75133) (isa
property75047 Property)
(adjacentTo property75047 property75133)
      (isa      group-of-plaintiff74903      (SetOfTypeFn      SocialBeing))
(possessiveRelation group-of-defendant75101 property75133)
      (isa      group-of-defendant75101      (SetOfTypeFn      SocialBeing))
(possessiveRelation group-of-plaintiff74903 property75047)

;;;The defendants began to build an addition.
(doneBy build75255 group-of-defendant75101)
(doneBy begin75234 group-of-defendant75101)
(isa addition75315 NewArtifact)
(isa  build75255  MakingSomething)  (situationBeginning  begin75234  build75255)
(outputsCreated build75255 addition75315)
(isa begin75234 BeginningAnActivity) (isa addition75315 RoomInAConstruction)

;;;The plaintiffs' gutters were over the defendants' property.
(over-UnderspecifiedLocation group-of-gutter78692 property75133)
(possessiveRelation group-of-plaintiff74903 group-of-gutter78692)
(isa group-of-gutter78692 (SetOfTypeFn GutterOnARoof))

;;;The gutters had existed for twenty years.
(objectActedOn exist74910 group-of-gutter78692) (doneBy exist74910 group-of-
gutter78692)
(measure exist74910 (YearsDuration 20))
```

```
(isa exist74910 SomethingExisting)
(isa gutter75559 GutterOnARoof)

;;;The defendants told the plaintiffs to remove the gutters.
(implies (and (isa tell75728 Requesting-CommunicationAct)
                        (recipientOfInfo tell75728 group-of-plaintiff74903)
                        (performedBy tell75728 group-of-defendant75101)
                        (situationTopic tell75728 remove75816))
            (and (isa remove75816 RemovingSomething)
                  (doneBy remove75816 group-of-plaintiff74903)
                  (objectRemoved remove75816 group-of-gutter78692)))

;;;The defendants built a new foundation on the plaintiffs' property.
(isa plaintiff74903 SocialBeing) (relationExistsInstance plaintiffs CourtCase
plaintiff74903)
(elementOf plaintiff74903 group-of-plaintiff74903)
(isa defendant75101 SocialBeing) (relationExistsInstance defendants CourtCase
defendant75101)
(elementOf defendant75101 group-of-defendant75101)
(isa group-of-defendant75101 (SetOfTypeFn SocialBeing))
(isa group-of-plaintiff74903 (SetOfTypeFn SocialBeing)) (isa build75255
BuildingSomething-Supervising)
(doneBy build75255 group-of-defendant75101) (outputsCreated build75255
foundation6927)
(eventOccursAt build75255 property75047)
(causes-PropProp
(and (isa build75255 BuildingSomething-Supervising) (doneBy build75255 group-
of-defendant75101))
(objectFoundInLocation foundation6927 property75047))
(owns group-of-plaintiff74903 property75047) (isa property75047 Property)
(isa foundation6927 FoundationOfABuilding) (possessiveRelation group-of-
plaintiff74903 property75047)

;;;The plaintiffs notified the defendants of this.
(isa notify76364 Informing)
(infoTransferred notify76364 (on-UnderspecifiedSurface foundation76126
property75047))
(recipientOfInfo notify76364 group-of-defendant75101)
(performedBy notify76364 group-of-plaintiff74903)

;;;The plaintiffs asked the defendants to stop building the addition.
;;; fix the "stopping" thing
(implies (and (isa ask76592 Requesting-CommunicationAct)
                  (recipientOfInfo ask76592 group-of-defendant75101)
                  (situationTopic ask76592 stop76682)
                  (performedBy ask76592 group-of-plaintiff74903))
            (and (temporalThingTerminated stop76682 build76078)
                  (doneBy build76078 group-of-defendant75101)
                  (doneBy stop76682 group-of-defendant75101)
                  (isa stop76682 DiscontinuingAnActivity)))

;;;The defendants destroyed the plaintiffs' gutters.
(doneBy destroy78828 group-of-defendant75101) (isa destroy78828
DestructionEvent)
(objectActedOn destroy78828 group-of-gutter78692)

;;;The defendants built the addition.
```

```
(doneBy build75255 group-of-defendant75101)
(isa addition75315 NewArtifact)
(isa build75255 MakingSomething)
(outputsCreated build75255 addition75315)
(isa addition75315 RoomInAConstruction)

;;;The defendants installed a drainage system.
(doneBy install78111 group-of-defendant75101)
(isa drainage-system78178 DrainageSystem)
(objectActedOn install78111 drainage-system78178)
(isa install78111 InstallingAnObject)

;;;The drainage system did not function.
(not (isa drainage-system79564 FunctionalObject))

;;;The drainage system damaged the plaintiffs' building.
(isa damage78543 IncurringPhysicalDamage) (objectActedOn damage78543
house74952)
      (doneBy damage78543 drainage-system78178)


;;;;;;;;;;;;;;;;;;;
(in-microtheory (cleanLanguageOutputMtFn-NoDiscourseNF Ariola_v_Nigro))
(genlMt (cleanLanguageOutputMtFn-NoDiscourseNF Ariola_v_Nigro)
(CaseLawCorpusMtFn Ariola_v_Nigro))

;;;The plaintiffs own a house.
(doneBy own74910 group-of-plaintiff74903)
      (isa group-of-plaintiff74903 (SetOfTypeFn SocialBeing)) (objectActedOn
own74910 house74952)
      (isa own74910 OwningSomething) (isa house74952 House-Modern) (elementOf
plaintiff74903 group-of-plaintiff74903)
      (relationExistsInstance plaintiffs CourtCase plaintiff74903) (isa
plaintiff74903 SocialBeing)

;;;The plaintiffs' property is adjacent to the defendants' property.
;;; add the plaintiffs own!
(isa property75133 Property) (owns group-of-defendant75101 property75133) (isa
property75047 Property)
(adjacentTo property75047 property75133)
      (isa group-of-plaintiff74903 (SetOfTypeFn SocialBeing))
(possessiveRelation group-of-defendant75101 property75133)
      (isa group-of-defendant75101 (SetOfTypeFn SocialBeing))
(possessiveRelation group-of-plaintiff74903 property75047)

;;;The defendants began to build an addition.
(doneBy build75255 group-of-defendant75101)
(doneBy begin75234 group-of-defendant75101)
(isa addition75315 NewArtifact)
(isa build75255 MakingSomething) (situationBeginning begin75234 build75255)
(outputsCreated build75255 addition75315)
(isa begin75234 BeginningAnActivity) (isa addition75315 RoomInAConstruction)

;;;The plaintiffs' gutters were over the defendants' property.
(over-UnderspecifiedLocation group-of-gutter78692 property75133)
(possessiveRelation group-of-plaintiff74903 group-of-gutter78692)
(isa group-of-gutter78692 (SetOfTypeFn GutterOnARoof))
```

```
;;;The gutters had existed for twenty years.
(objectActedOn exist74910 group-of-gutter78692) (doneBy exist74910 group-of-gutter78692)
(measure exist74910 (YearsDuration 20))
(isa exist74910 SomethingExisting)
(isa gutter75559 GutterOnARoof)

;;;The defendants told the plaintiffs to remove the gutters.
(implies (and (isa tell75728 Requesting-CommunicationAct)
                        (recipientOfInfo tell75728 group-of-plaintiff74903)
                        (performedBy tell75728 group-of-defendant75101)
                        (situationTopic tell75728 remove75816))
            (and (isa remove75816 RemovingSomething)
                  (doneBy remove75816 group-of-plaintiff74903)
                  (objectRemoved remove75816 group-of-gutter78692)))

;;;The defendants built a new foundation on the plaintiffs' property.
(isa plaintiff74903 SocialBeing) (relationExistsInstance plaintiffs CourtCase plaintiff74903)
(elementOf plaintiff74903 group-of-plaintiff74903)
(isa defendant75101 SocialBeing) (relationExistsInstance defendants CourtCase defendant75101)
(elementOf defendant75101 group-of-defendant75101)
(isa group-of-defendant75101 (SetOfTypeFn SocialBeing))
(isa group-of-plaintiff74903 (SetOfTypeFn SocialBeing)) (isa build75255 BuildingSomething-Supervising)
(doneBy build75255 group-of-defendant75101) (outputsCreated build75255 foundation6927)
(eventOccursAt build75255 property75047)
(causes-PropProp
(and (isa build75255 BuildingSomething-Supervising) (doneBy build75255 group-of-defendant75101))
(objectFoundInLocation foundation6927 property75047))
(owns group-of-plaintiff74903 property75047) (isa property75047 Property)
(isa foundation6927 FoundationOfABuilding) (possessiveRelation group-of-plaintiff74903 property75047)

;;;The plaintiffs notified the defendants of this.
(isa notify76364 Informing)
(infoTransferred notify76364 (on-UnderspecifiedSurface foundation76126 property75047))
(recipientOfInfo notify76364 group-of-defendant75101)
(performedBy notify76364 group-of-plaintiff74903)

;;;The plaintiffs asked the defendants to stop building the addition.
;;; fix the "stopping" thing
(implies (and (isa ask76592 Requesting-CommunicationAct)
                  (recipientOfInfo ask76592 group-of-defendant75101)
                  (situationTopic ask76592 stop76682)
                  (performedBy ask76592 group-of-plaintiff74903))
            (and (temporalThingTerminated stop76682 build76078)
                  (doneBy build76078 group-of-defendant75101)
                  (doneBy stop76682 group-of-defendant75101)
                  (isa stop76682 DiscontinuingAnActivity)))

;;;The defendants destroyed the plaintiffs' gutters.
```

```
(doneBy destroy78828 group-of-defendant75101) (isa destroy78828 DestructionEvent)
(objectActedOn destroy78828 group-of-gutter78692)

;;;The defendants built the addition.
(doneBy build75255 group-of-defendant75101)
(isa addition75315 NewArtifact)
(isa build75255 MakingSomething)
(outputsCreated build75255 addition75315)
(isa addition75315 RoomInAConstruction)

;;;The defendants installed a drainage system.
(doneBy install78111 group-of-defendant75101)
(isa drainage-system78178 DrainageSystem)
(objectActedOn install78111 drainage-system78178)
(isa install78111 InstallingAnObject)

;;;The drainage system did not function.
(not (isa drainage-system79564 FunctionalObject))

;;;The drainage system damaged the plaintiffs' building.
(isa damage78543 IncurringPhysicalDamage) (objectActedOn damage78543 house74952)
      (doneBy damage78543 drainage-system78178)

;;;;;;;;;;;;;;;;;;;;;

(in-microtheory (LegalCaseMtFn Ariola_v_Nigro))
(genlMt (LegalCaseMtFn Ariola_v_Nigro) (CaseLawCorpusMtFn Ariola_v_Nigro))

(isa Ariola_v_Nigro CourtCase)
(plaintiffs Ariola_v_Nigro group-of-plaintiff74903)
(defendants Ariola_v_Nigro group-of-defendant75101)

;;;The plaintiffs own a house.
(doneBy own74910 group-of-plaintiff74903)
      (isa group-of-plaintiff74903 (SetOfTypeFn Plaintiff)) (objectActedOn own74910 house74952)
      (isa own74910 OwningSomething) (isa house74952 House-Modern) (elementOf plaintiff74903 group-of-plaintiff74903)
      (isa plaintiff74903 Plaintiff)

;;;The plaintiffs' property is adjacent to the defendants' property.
(isa property75133 Property) (owns group-of-defendant75101 property75133) (isa property75047 Property)
(adjacentTo property75047 property75133)
(isa group-of-defendant75101 (SetOfTypeFn Defendant))
(owns group-of-plaintiff74903 property75047)
(possessiveRelation group-of-defendant75101 property75133)
(possessiveRelation group-of-plaintiff74903 property75047)

;;;The defendants began to build an addition.
(doneBy build75255 group-of-defendant75101)
(doneBy begin75234 group-of-defendant75101)
(isa addition75315 NewArtifact)
(isa build75255 MakingSomething) (situationBeginning begin75234 build75255)
(outputsCreated build75255 addition75315)
```

```
(isa begin75234 BeginningAnActivity) (isa addition75315 RoomInAConstruction)

;;;The plaintiffs' gutters were over the defendants' property.
(over-UnderspecifiedLocation group-of-gutter78692 property75133)
(possessiveRelation group-of-plaintiff74903 group-of-gutter78692)
(isa group-of-gutter78692 (SetOfTypeFn GutterOnARoof))

;;;The gutters had existed for twenty years.
(objectActedOn exist74910 group-of-gutter78692) (doneBy exist74910 group-of-
gutter78692)
(measure exist74910 (YearsDuration 20))
(isa exist74910 SomethingExisting)
(isa gutter75559 GutterOnARoof)

;;;The defendants told the plaintiffs to remove the gutters.
(implies (and (isa tell75728 Requesting-CommunicationAct)
                        (recipientOfInfo tell75728 group-of-plaintiff74903)
                        (performedBy tell75728 group-of-defendant75101)
                        (situationTopic tell75728 remove75816))
            (and (isa remove75816 RemovingSomething)
                  (doneBy remove75816 group-of-plaintiff74903)
                  (objectRemoved remove75816 group-of-gutter78692)))

;;;The defendants built a new foundation on the plaintiffs' property.
(isa build75255 BuildingSomething-Supervising)
(doneBy build75255 group-of-defendant75101)
(outputsCreated build75255 foundation6927)
(eventOccursAt build75255 property75047)
(causes-PropProp
      (and (isa build75255 BuildingSomething-Supervising) (doneBy build75255
group-of-defendant75101))
      (objectFoundInLocation foundation6927 property75047))
(isa foundation6927 FoundationOfABuilding)

;;;The plaintiffs notified the defendants of this.
(isa notify76364 Informing)
(infoTransferred notify76364 (on-UnderspecifiedSurface foundation76126
property75047))
(recipientOfInfo notify76364 group-of-defendant75101)
(performedBy notify76364 group-of-plaintiff74903)

;;;The plaintiffs asked the defendants to stop building the addition.
(implies (and (isa ask76592 Requesting-CommunicationAct)
                  (recipientOfInfo ask76592 group-of-defendant75101)
                  (situationTopic ask76592 stop76682)
                  (performedBy ask76592 group-of-plaintiff74903))
            (and (temporalThingTerminated stop76682 build75255)
                  (doneBy stop76682 group-of-defendant75101)
                  (isa stop76682 DiscontinuingAnActivity)))

;;;The defendants destroyed the plaintiffs' gutters.
(doneBy destroy78828 group-of-defendant75101) (isa destroy78828
DestructionEvent)
(objectActedOn destroy78828 group-of-gutter78692)

;;;The defendants built the addition.
(doneBy build75255 group-of-defendant75101)
```

```
(isa addition75315 NewArtifact)
(isa build75255 MakingSomething)
(outputsCreated build75255 addition75315)
(isa addition75315 RoomInAConstruction)

;;;The defendants installed a drainage system.
(doneBy install78111 group-of-defendant75101)
(isa drainage-system78178 DrainageSystem)
(objectActedOn install78111 drainage-system78178)
(isa install78111 InstallingAnObject)

;;;The drainage system did not function.
(not (isa drainage-system79564 FunctionalObject))

;;;The drainage system damaged the plaintiffs' building.
(isa   damage78543   IncurringPhysicalDamage)   (objectActedOn   damage78543
house74952)
      (doneBy damage78543 drainage-system78178)

(in-microtheory (LegalCaseConclusionMtFn Ariola_v_Nigro))
(genlMt      (LegalCaseConclusionMtFn      Ariola_v_Nigro)      (LegalCaseMtFn
Ariola_v_Nigro))

(trespassOnPropertyByAction group-of-defendant75101 property75047 build75255)
(hasEasementOnProperty group-of-plaintiff74903 property75133)

(in-microtheory (LegalCaseConclusion-NegatedMtFn Ariola_v_Nigro))
(not    (trespassOnPropertyByAction   group-of-defendant75101   property75047
build75255))
(not (hasEasementOnProperty group-of-plaintiff74903 property75133))

(in-microtheory (LegalCaseConclusion-ReversedMtFn Ariola_v_Nigro))
(trespassOnPropertyByAction group-of-plaintiff74903 property75047 build75255)
(hasEasementOnProperty group-of-defendant75101 property75133)

(in-microtheory (LegalCaseConclusion-NegatedReversedMtFn Ariola_v_Nigro))
(not    (trespassOnPropertyByAction   group-of-plaintiff74903   property75047
build75255))
(not (hasEasementOnProperty group-of-defendant75101 property75133))
```

```
(in-microtheory CaseLawCorpusMt)

(isa Ayoubi_v_Dart CaseLawCorpusCase)

(in-microtheory (CaseLawCorpusMtFn Ayoubi_v_Dart))

(caseType Ayoubi_v_Dart Assault)
(caseType Ayoubi_v_Dart Battery)

(caseValence Ayoubi_v_Dart Assault Negative)
(caseValence Ayoubi_v_Dart Battery Negative)

(caseName Ayoubi_v_Dart "Ayoubi v. Dart")

(caseCourt Ayoubi_v_Dart "United States District Court for the Northern District
of Illinois")

(caseYear Ayoubi_v_Dart "2015")

(caseReporter Ayoubi_v_Dart "2015 U.S. Dist. LEXIS 33446")

(caseOriginalText Ayoubi_v_Dart "Plaintiff Firas Ayoubi, a pretrial detainee at
the Cook County Jail, filed this civil rights action against Cook County Sheriff
Tom Dart, Cook County Department of Corrections Executive Director Murphy,
Superintendents Menella, Queen, and Everheart, and Cermak Director Dr. Kahn. In
December of 2012 inmates with contagious diseases, who had been labeled
'quarantine/isolation,' were housed in Division 5 along with Plaintiff and other
inmates. Plaintiff states that the Defendants were aware that contagious inmates
were being housed with non-contagious inmates; that he was not moved despite
his requests; and that he contracted an illness, which caused high fever and
extensive coughing for two weeks.")

(caseConclusion Ayoubi_v_Dart "Defendant did not commit an assault.")

(caseConclusion Ayoubi_v_Dart "Defendant did not commit a battery.")

(caseSimpleText Ayoubi_v_Dart "The plaintiff is imprisoned in jail. The
defendant controls the jail. The defendant placed sick inmates in the
plaintiff's cell. The defendant knew the sick inmates were contagious. The
plaintiff asked the defendant to move the plaintiff to a different cell. The
plaintiff was not moved. The plaintiff became sick.")

(in-microtheory (rawLanguageOutputMtFn Ayoubi_v_Dart))
(genlMt (rawLanguageOutputMtFn Ayoubi_v_Dart) (CaseLawCorpusMtFn
Ayoubi_v_Dart))

(and (eventOccursAt imprison108748 jail108862) (objectFoundInLocation
plaintiff108732 jail108862)
      (isa imprison108748 Imprisoning) (captive imprison108748 plaintiff108732)
(isa jail108862 Prison)
      (isa plaintiff108732 SocialBeing) (relationExistsInstance plaintiffs
CourtCase plaintiff108732))
(and (isa jail109048 Prison) (performedBy control109019 defendant109007)
      (isa control109019 ExercisingAuthoritativeControlOverSomething)
(objectControlled control109019 jail109048)
      (isa defendant109007 SocialBeing) (relationExistsInstance defendants
CourtCase defendant109007))
```

```
(and   (doneBy   place109134   defendant109007)   (objectPlaced   place109134
inmate109176)
      (isa place109134 PuttingSomethingSomewhere) (stateOfHealth inmate109176
Sick)
        (isa cell109271 PrisonCell) (near inmate109176 plaintiff108732)
        (causes-PropProp
            (and (isa place109134 PuttingSomethingSomewhere)
                  (doneBy place109134 defendant109007) (to-Generic place109134
cell109271))
            (objectFoundInLocation inmate109176 cell109271))
      (possessiveRelation plaintiff108732 cell109271) (to-Generic place109134
cell109271)
      (relationExistsInstance   plaintiffs   CourtCase   plaintiff108732)   (isa
defendant109007 SocialBeing)
      (isa  inmate109176  Prisoner-Legal)  (relationExistsInstance  defendants
CourtCase defendant109007)
      (isa plaintiff108732 SocialBeing))
(and (implies
       (and    (isa    know108304    (HavingPropositionalAttitudeFn    knows))
(subjectOfMentalSituation know108304 defendant109007))
       (and (by-Underspecified be108351 group-of-inmate108336)
            (hasCommunicabilityLevel                      group-of-inmate108336
HighLevelOfCommunicability)
            (relationInstanceExists by-Underspecified be108351
             (CollectionSubsetFn Prisoner-Legal
              (TheSetOf ?inmate108336 (and (isa ?inmate108336 Prisoner-Legal)
(stateOfHealth ?inmate108336 Sick)))))
            (relationExistsInstance hasCommunicabilityLevel
             (CollectionSubsetFn Prisoner-Legal
              (TheSetOf ?inmate108336 (and (isa ?inmate108336 Prisoner-Legal)
(stateOfHealth ?inmate108336 Sick))))
             HighLevelOfCommunicability)
            (isa be108351 Situation) (isa group-of-inmate108336 (SetOfTypeFn
SocialBeing))
            (elementOf inmate108336 group-of-inmate108336) (isa inmate108336
Prisoner-Legal)
            (stateOfHealth inmate108336 Sick)))
      (isa  defendant109007  SocialBeing)  (relationExistsInstance  defendants
CourtCase defendant109007))
(and  (recipientOfInfo  ask109616  defendant109007)  (to-Generic  move109701
cell109887) (isa cell109887 PrisonCell)
      (situationTopic  ask109616  move109701)  (isa  move109701  (CausingFn
MovementEvent)) (doneBy move109701 defendant109007)
      (isa  ask109616  Requesting-CommunicationAct)  (objectActedOn  move109701
inmate108336)
      (relationExistsInstance defendants CourtCase defendant109007)
      (relationExistsInstance   plaintiffs   CourtCase   inmate108336)   (isa
inmate108336 SocialBeing)
      (isa defendant109007 SocialBeing) (performedBy ask109616 inmate108336))
(and (not (and (isa move108610 (CausingFn MovementEvent)) (doneBy move108610
inmate108336)))
      (relationExistsInstance   plaintiffs   CourtCase   inmate108336)   (isa
inmate108336 SocialBeing))
(and  (stateOfHealth  inmate108336  Sick)  (relationExistsInstance  plaintiffs
CourtCase inmate108336)
      (isa inmate108336 SocialBeing))
```

```
(in-microtheory (cleanLanguageOutputMtFn Ayoubi_v_Dart))
(genlMt     (cleanLanguageOutputMtFn     Ayoubi_v_Dart)     (CaseLawCorpusMtFn
Ayoubi_v_Dart))

;;;The plaintiff is imprisoned in jail.
(eventOccursAt     imprison108748     jail108862)     (objectFoundInLocation
plaintiff108732 jail108862)
      (isa imprison108748 Imprisoning) (captive imprison108748 plaintiff108732)
(isa jail108862 Prison)
      (isa  plaintiff108732  SocialBeing)  (relationExistsInstance  plaintiffs
CourtCase plaintiff108732)

;;;The defendant controls the jail.
(isa jail108862 Prison) (performedBy control109019 defendant109007)
      (isa     control109019     ExercisingAuthoritativeControlOverSomething)
(objectControlled control109019 jail108862)
      (isa  defendant109007  SocialBeing)  (relationExistsInstance  defendants
CourtCase defendant109007)

;;;The defendant placed sick inmates in the plaintiff's cell.
(doneBy  place109134  defendant109007)  (objectPlaced  place109134  group-of-
inmate108336)
      (isa  place109134  PuttingSomethingSomewhere)  (stateOfHealth  group-of-
inmate108336 Sick)
        (isa   cell109271   PrisonCell)   (isa   cell109271   PrisonCell)   (near
inmate109176 plaintiff108732)
        (causes-PropProp
            (and (isa place109134 PuttingSomethingSomewhere)
                  (doneBy place109134 defendant109007) (to-Generic place109134
cell109271))
            (objectFoundInLocation inmate109176 cell109271))
      (possessiveRelation plaintiff108732 cell109271) (to-Generic place109134
cell109271)
      (elementOf inmate108336 group-of-inmate108336)
        (isa inmate108336 Prisoner-Legal)

;;;The defendant knew the sick inmates were contagious.
(implies   (and   (isa   know108304   (HavingPropositionalAttitudeFn   knows))
(subjectOfMentalSituation know108304 defendant109007))
       (and           (hasCommunicabilityLevel           group-of-inmate108336
HighLevelOfCommunicability)
            (stateOfHealth group-of-inmate108336 Sick)))

;;;The plaintiff asked the defendant to move the plaintiff to a different cell.
(implies (and (recipientOfInfo ask109616 defendant109007)
                        (situationTopic ask109616 move109701)
                        (isa ask109616 Requesting-CommunicationAct)
                        (performedBy ask109616 plaintiff108732))
            (and (to-Generic move109701 cell109887)
                  (isa cell109887 PrisonCell)
                  (isa move109701 (CausingFn MovementEvent))
                  (doneBy move109701 defendant109007)
                  (objectActedOn move109701 plaintiff108732)))

;;;The plaintiff was not moved.
(not (and (to-Generic move109701 cell109887)
                  (isa cell109887 PrisonCell)
```

```
                (isa move109701 (CausingFn MovementEvent))
                (doneBy move109701 defendant109007)
                (objectActedOn move109701 plaintiff108732)))

;;;The plaintiff became sick.
(stateOfHealth plaintiff108732 Sick)

(in-microtheory (LegalCaseMtFn Ayoubi_v_Dart))
(genlMt (LegalCaseMtFn Ayoubi_v_Dart) (CaseLawCorpusMtFn Ayoubi_v_Dart))

(isa Ayoubi_v_Dart CourtCase)
(plaintiffs Ayoubi_v_Dart plaintiff108732)
(defendants Ayoubi_v_Dart defendant109007)

;;;The plaintiff is imprisoned in jail.
(eventOccursAt     imprison108748     jail108862)     (objectFoundInLocation
plaintiff108732 jail108862)
      (isa imprison108748 Imprisoning) (captive imprison108748 plaintiff108732)
(isa jail108862 Prison)
      (isa plaintiff108732 Plaintiff)

;;;The defendant controls the jail.
(isa jail108862 Prison) (performedBy control109019 defendant109007)
      (isa     control109019     ExercisingAuthoritativeControlOverSomething)
(objectControlled control109019 jail108862)
      (isa defendant109007 Defendant)

;;;The defendant placed sick inmates in the plaintiff's cell.
(doneBy  place109134  defendant109007)  (objectPlaced  place109134  group-of-
inmate108336)
      (isa  place109134  PuttingSomethingSomewhere)  (stateOfHealth  group-of-
inmate108336 Sick)
        (isa   cell109271   PrisonCell)   (isa   cell109271   PrisonCell)   (near
inmate109176 plaintiff108732)
        (causes-PropProp
            (and (isa place109134 PuttingSomethingSomewhere)
                  (doneBy place109134 defendant109007) (to-Generic place109134
cell109271))
            (objectFoundInLocation inmate109176 cell109271))
      (possessiveRelation plaintiff108732 cell109271) (to-Generic place109134
cell109271)
      (elementOf inmate108336 group-of-inmate108336)
        (isa inmate108336 Prisoner-Legal)

;;;The defendant knew the sick inmates were contagious.
(implies   (and   (isa   know108304   (HavingPropositionalAttitudeFn   knows))
(subjectOfMentalSituation know108304 defendant109007))
       (and           (hasCommunicabilityLevel           group-of-inmate108336
HighLevelOfCommunicability)
            (stateOfHealth group-of-inmate108336 Sick)))

;;;The plaintiff asked the defendant to move the plaintiff to a different cell.
(implies (and (recipientOfInfo ask109616 defendant109007)
                        (situationTopic ask109616 move109701)
                        (isa ask109616 Requesting-CommunicationAct)
                        (performedBy ask109616 plaintiff108732))
            (and (to-Generic move109701 cell109887)
```

```
                  (isa cell109887 PrisonCell)
                  (isa move109701 (CausingFn MovementEvent))
                  (doneBy move109701 defendant109007)
                  (objectActedOn move109701 plaintiff108732)))

;;;The plaintiff was not moved.
(not (and (to-Generic move109701 cell109887)
                  (isa cell109887 PrisonCell)
                  (isa move109701 (CausingFn MovementEvent))
                  (doneBy move109701 defendant109007)
                  (objectActedOn move109701 plaintiff108732)))

;;;The plaintiff became sick.
;;;  adding  a  big'un  here;  the  original  statement  was  just  (stateOfHealth
inmate108336 Sick)
(causes-PropProp (and (near inmate108336 plaintiff108732)
                                    (stateOfHealth group-of-inmate108336 Sick)
                                    (hasCommunicabilityLevel         group-of-
inmate108336 HighLevelOfCommunicability))
                        (stateOfHealth inmate108336 Sick))

(in-microtheory (LegalCaseConclusionMtFn Ayoubi_v_Dart))
(genlMt (LegalCaseConclusionMtFn Ayoubi_v_Dart) (LegalCaseMtFn Ayoubi_v_Dart))

(not (assaultsPartyByDoing defendant109007 plaintiff108732 place109134))
(not (battersPartyByDoing defendant109007 plaintiff108732 place109134))

(in-microtheory (LegalCaseConclusion-NegatedMtFn Ayoubi_v_Dart))
(assaultsPartyByDoing defendant109007 plaintiff108732 place109134)
(battersPartyByDoing defendant109007 plaintiff108732 place109134)

(in-microtheory (LegalCaseConclusion-ReversedMtFn Ayoubi_v_Dart))
(not (assaultsPartyByDoing plaintiff108732 defendant109007 place109134))
(not (battersPartyByDoing plaintiff108732 defendant109007 place109134))

(in-microtheory (LegalCaseConclusion-NegatedReversedMtFn Ayoubi_v_Dart))
(assaultsPartyByDoing plaintiff108732 defendant109007 place109134)
(battersPartyByDoing plaintiff108732 defendant109007 place109134)
```

(in-microtheory CaseLawCorpusMt)

(isa Banks_v_Chicago CaseLawCorpusCase)

(in-microtheory (CaseLawCorpusMtFn Banks_v_Chicago))

(caseType Banks_v_Chicago Assault)
(caseType Banks_v_Chicago Self_Defense)

(caseValence Banks_v_Chicago Assault Positive)
(caseValence Banks_v_Chicago Self_Defense Positive)

(caseName Banks_v_Chicago "Banks v. Chicago")

(caseCourt Banks_v_Chicago "Appellate Court of Illinois, First District, Fifth Division")

(caseYear Banks_v_Chicago "1973")

(caseReporter Banks_v_Chicago "11 Ill. App. 3d 543")

(caseOriginalText Banks_v_Chicago "On November 24, 1960, at about 2:00 or 2:30 A.M., an off-duty Chicago police officer, Floyd Pace, was in Sam's Chicken Shack. This establishment housed both a restaurant and a tavern, the two being connected by a corridor about three feet wide and eight to ten feet long. Pace had been drinking a beer in the tavern for about 15 minutes when he got up from the bar and walked to the connecting corridor. He was midway through the corridor when he saw plaintiff's decedent, Otis Banks, approaching him from the opposite direction. He observed that Banks could not support himself, so he backed up against the wall to let him by. As Banks passed Pace in the hallway, he bumped into the officer and said, 'Get the fuck out of my way.' The officer took out his star and identified himself as a policeman in order to calm down Banks. Then Pace continued walking toward the restaurant while Banks Walked in the opposite direction. Pace was walking past the first booth off the corridor in the restaurant when Banks ran around in front of him and pushed the barrel of a .25-caliber automatic revolver into his lower lip. Pace's mouth began to bleed, he saw a flash, and he thought that Banks had pulled the trigger. Banks ran toward the south wall of the restaurant, and Pace called out, 'Halt, Police.' When Banks reached the wall, he turned around, put both hands on the gun, and pointed it toward the officer. Pace pulled out his own gun and fired it once at Banks. Pace had been carrying his service revolver on his person, but had not exhibited it in the establishment prior to this time. Banks slumped to the floor in a sitting position with his gun in his lap. He raised the gun a second time and aimed directly at Pace. By the motion of Banks' hand, it appeared to Pace that Banks was attempting to pull the trigger. Pace fired a second time, and Banks collapsed on the floor. The restaurant was filled almost to capacity, and at the time Banks aimed the gun at Pace, the counter, about six feet on Pace's left, was crowded, although there were no tables nearby.")

(caseConclusion Banks_v_Chicago "The defendant acted in self-defense.")
(caseConclusion Banks_v_Chicago "The plaintiff assaulted the defendant.")

(caseSimpleText Banks_v_Chicago "The defendant was a police officer. The defendant was in a restaurant. The plaintiff was intoxicated. The defendant told the plaintiff the defendant was a police officer. Then the plaintiff pointed a gun at the defendant. Then the plaintiff ran away. The defendant told the plaintiff to stop running. Then the plaintiff pointed a gun at the defendant again. Then the defendant fired the gun at the plaintiff. Then the plaintiff aimed the gun at the defendant again. Then the defendant shot the plaintiff again.")

```
(in-microtheory (rawLanguageOutputMtFn Banks_v_Chicago))
(genlMt (rawLanguageOutputMtFn Banks_v_Chicago) (CaseLawCorpusMtFn Banks_v_Chicago))

((and (relationExistsInstance defendants CourtCase defendant107259) (isa defendant107259 PoliceOfficer-Municipal) (isa defendant107259 SocialBeing))
 (and (objectFoundInLocation defendant107259 restaurant107428) (isa restaurant107428 RestaurantSpace) (by-Underspecified be107386 defendant107259)
      (isa defendant107259 SocialBeing) (relationExistsInstance defendants CourtCase defendant107259) (isa be107386 Situation))
 (and (by-Underspecified be107477 police-officer107344) (isa be107477 Situation) (hasState police-officer107344 Intoxicated)
      (relationExistsInstance plaintiffs CourtCase police-officer107344) (isa police-officer107344 SocialBeing))
 (and (and (isa point107946 MakingAThreat) (threateningAgent point107946 police-officer107344) (threatenedAgent point107946 defendant107259)) (isa gun107990 Weapon)
      (doneBy point107946 police-officer107344) (intendedToLocation point107946 defendant107259) (instrument-Generic point107946 gun107990)
      (isa point107946 AimingSomething) (relationExistsInstance nextInSequence List point107946) (doneBy gun107990 police-officer107344)
      (isa police-officer107344 SocialBeing) (relationExistsInstance defendants CourtCase defendant107259)
      (relationExistsInstance plaintiffs CourtCase police-officer107344) (isa defendant107259 SocialBeing) (isa gun107990 Gun))
 (and (isa run108201 Running) (performedBy run108201 police-officer107344) (isa police-officer107344 SocialBeing)
      (relationExistsInstance plaintiffs CourtCase police-officer107344) (relationExistsInstance nextInSequence List run108201))
 (and (isa tell108282 Requesting-CommunicationAct) (doneBy stop108382 police-officer107344) (isa run108489 Running) (temporalThingTerminated stop108382 run108489)
      (performedBy run108489 police-officer107344) (performedBy tell108282 defendant107259) (isa stop108382 DiscontinuingAnActivity)
      (situationTopic tell108282 stop108382) (recipientOfInfo tell108282 police-officer107344) (purposeOf-Generic run108489 stop108382)
      (relationExistsInstance defendants CourtCase defendant107259) (relationExistsInstance plaintiffs CourtCase police-officer107344)
      (isa police-officer107344 SocialBeing) (isa defendant107259 SocialBeing))
 (and (and (isa point108778 MakingAThreat) (threateningAgent point108778 police-officer107344) (threatenedAgent point108778 defendant107259)) (isa gun108822 Weapon)
      (doneBy point108778 police-officer107344) (intendedToLocation point108778 defendant107259) (instrument-Generic point108778 gun108822)
      (isa point108778 AimingSomething) (isa police-officer107344 SocialBeing) (isa gun108822 Gun) (doneBy gun108822 police-officer107344)
      (relationExistsInstance nextInSequence List point108778) (isa defendant107259 SocialBeing) (isa point108778 RepeatedEventType)
      (relationExistsInstance defendants CourtCase defendant107259) (relationExistsInstance plaintiffs CourtCase police-officer107344))
 (and (isa gun107990 Weapon) (doneBy fire109059 defendant107259) (to-Generic fire109059 police-officer107344) (launcherInShooting fire109059 gun107990)
```

```
      (isa fire109059 ShootingAGun) (isa gun107990 Gun) (isa defendant107259
SocialBeing) (doneBy gun107990 defendant107259) (isa police-officer107344
SocialBeing)
      (relationExistsInstance nextInSequence List fire109059)
(relationExistsInstance plaintiffs CourtCase police-officer107344)
      (relationExistsInstance defendants CourtCase defendant107259))
 (and (and (isa point108778 MakingAThreat) (threateningAgent point108778
police-officer107344) (threatenedAgent point108778 defendant107259)) (isa
gun107990 Weapon)
      (doneBy point108778 police-officer107344) (intendedToLocation
point108778 defendant107259) (instrument-Generic point108778 gun107990)
      (isa point108778 AimingSomething) (isa police-officer107344 SocialBeing)
(isa defendant107259 SocialBeing) (doneBy gun107990 police-officer107344)
      (isa point108778 RepeatedEventType) (isa gun107990 Gun)
(relationExistsInstance defendants CourtCase defendant107259)
      (relationExistsInstance plaintiffs CourtCase police-officer107344)
(relationExistsInstance nextInSequence List point108778))
 (and (intendedToLocation shoot109629 police-officer107344) (isa shoot109629
ShootingAndHittingSomething) (doneBy shoot109629 defendant107259)
      (isa defendant107259 SocialBeing) (isa shoot109629 RepeatedEventType)
(relationExistsInstance plaintiffs CourtCase police-officer107344)
      (relationExistsInstance defendants CourtCase defendant107259)
(relationExistsInstance nextInSequence List shoot109629) (isa police-
officer107344 SocialBeing))
 (and (recipientOfInfo be107386 police-officer107344) (performedBy be107386
defendant107259) (isa be107386 Informing) (isa defendant109905 PoliceOfficer-
Municipal)
      (relationExistsInstance defendants CourtCase defendant107259) (isa
defendant107259 SocialBeing) (isa police-officer107344 SocialBeing)
      (relationExistsInstance plaintiffs CourtCase police-officer107344)))

(in-microtheory (cleanLanguageOutputMtFn Banks_v_Chicago))
(genlMt (cleanLanguageOutputMtFn Banks_v_Chicago) (CaseLawCorpusMtFn
Banks_v_Chicago))

;;;The defendant was a police officer.
(relationExistsInstance defendants CourtCase defendant107259)
(isa defendant107259 PoliceOfficer-Municipal) (isa defendant107259 SocialBeing)

;;;The defendant was in a restaurant.
(objectFoundInLocation defendant107259 restaurant107428)
(isa restaurant107428 RestaurantSpace)
      (isa defendant107259 SocialBeing) (relationExistsInstance defendants
CourtCase defendant107259)

;;;The plaintiff was intoxicated.
(hasState plaintiff107470 Intoxicated)
(relationExistsInstance plaintiffs CourtCase plaintiff107470)
(isa plaintiff107470 SocialBeing)

;;;The defendant told the plaintiff the defendant was a police officer.
(recipientOfInfo tell110361 plaintiff107470)
(performedBy tell110361 defendant107259)
(isa tell110361 Informing)
(isa defendant109905 PoliceOfficer-Municipal)
(relationExistsInstance defendants CourtCase defendant107259) (isa
defendant107259 SocialBeing) (isa plaintiff107470 SocialBeing)
```

```
      (relationExistsInstance plaintiffs CourtCase plaintiff107470)

;;;Then the plaintiff pointed a gun at the defendant.
(relationExistsInstance nextInSequence List point107946)
(isa point107946 MakingAThreat)
(threateningAgent point107946 plaintiff107470)
(threatenedAgent point107946 defendant107259)
(isa gun107990 Weapon)
(doneBy point107946 plaintiff107470)
(intendedToLocation point107946 defendant107259)
(instrument-Generic point107946 gun107990)
(isa point107946 AimingSomething)
(isa plaintiff107470 SocialBeing)
(relationExistsInstance defendants CourtCase defendant107259)
(relationExistsInstance plaintiffs CourtCase plaintiff107470)
(isa defendant107259 SocialBeing) (isa gun107990 Gun)

;;;Then the plaintiff ran away.
(relationExistsInstance nextInSequence List run108489)
(isa run108489 Running) (performedBy run108489 plaintiff107470)
(isa plaintiff107470 SocialBeing)
(relationExistsInstance plaintiffs CourtCase plaintiff107470)

;;;The defendant told the plaintiff to stop running.
(isa tell108282 Requesting-CommunicationAct)
(doneBy stop108382 plaintiff107470) (isa run108489 Running)
(temporalThingTerminated stop108382 run108489)
      (performedBy run108489 plaintiff107470)
(performedBy      tell108282      defendant107259)      (isa      stop108382
DiscontinuingAnActivity)
      (situationTopic tell108282 stop108382)
        (recipientOfInfo tell108282 plaintiff107470)
      (relationExistsInstance defendants CourtCase defendant107259)
        (relationExistsInstance plaintiffs CourtCase run108489)
      (isa run108489 SocialBeing) (isa defendant107259 SocialBeing)

;;;Then the plaintiff pointed a gun at the defendant again.
(relationExistsInstance  nextInSequence  List  point108778)(isa  point108778
MakingAThreat)
(threateningAgent point108778 plaintiff107470)
(threatenedAgent point108778 defendant107259) (isa gun107990 Weapon)
(doneBy point108778 plaintiff107470)
(intendedToLocation point108778 defendant107259)
(instrument-Generic point108778 gun107990)
(isa point108778 AimingSomething)
(isa plaintiff107470 SocialBeing)
(isa gun107990 Gun)
(isa defendant107259 SocialBeing) (isa point108778 RepeatedEventType)
      (relationExistsInstance defendants CourtCase defendant107259)
        (relationExistsInstance plaintiffs CourtCase plaintiff107470)

;;;Then the defendant fired the gun at the plaintiff.
(relationExistsInstance nextInSequence List fire109059)
(isa gun107991 Weapon)
(doneBy fire109059 defendant107259)
(to-Generic fire109059 plaintiff107470)
(launcherInShooting fire109059 gun107991)
```

```
      (isa fire109059 ShootingAGun) (isa gun107991 Gun) (isa defendant107259
SocialBeing)  (isa plaintiff107470 SocialBeing)
       (relationExistsInstance plaintiffs CourtCase plaintiff107470)
      (relationExistsInstance defendants CourtCase defendant107259)
(startsAfterStartingOf fire109059 point108778)
(isa fire109059 AttackOnTangible)
(performedBy fire109059 defendant107259)
(objectAttacked fire109059 plaintiff107470)
(purposeInEvent    defendant107259    fire109059    (damages    fire109059
plaintiff107470))

;;;Then the plaintiff aimed the gun at the defendant again.
(relationExistsInstance nextInSequence List aim108778)
(isa  aim108778  MakingAThreat)  (threateningAgent  aim108778  plaintiff107470)
(threatenedAgent aim108778 defendant107259) (isa gun107990 Weapon)
      (doneBy   aim108778   plaintiff107470)   (intendedToLocation   aim108778
defendant107259) (instrument-Generic aim108778 gun107990)
      (isa aim108778 AimingSomething) (isa plaintiff107470 SocialBeing)
        (isa defendant107259 SocialBeing)
      (isa aim108778 RepeatedEventType) (isa gun107990 Gun)
        (relationExistsInstance defendants CourtCase defendant107259)
      (relationExistsInstance plaintiffs CourtCase plaintiff107470)

;;;Then the defendant shot the plaintiff again.
(relationExistsInstance nextInSequence List shoot109629)
(isa shoot109629 AttackOnTangible)
(performedBy shoot109629 defendant107259)
(objectAttacked shoot109629 plaintiff107470)
(damages shoot109629 plaintiff107470)
(intendedToLocation shoot109629 plaintiff107470)
(isa shoot109629 ShootingAndHittingSomething)
(doneBy shoot109629 defendant107259)
      (isa  defendant107259  SocialBeing)  (isa  shoot109629  RepeatedEventType)
(relationExistsInstance plaintiffs CourtCase plaintiff107470)
      (relationExistsInstance defendants CourtCase defendant107259)
        (isa plaintiff107470 SocialBeing)

(in-microtheory (cleanLanguageOutputMtFn-NoDiscourseNF Banks_v_Chicago))
(genlMt        (cleanLanguageOutputMtFn-NoDiscourseNF        Banks_v_Chicago)
(CaseLawCorpusMtFn Banks_v_Chicago))

;;;The defendant was a police officer.
(relationExistsInstance defendants CourtCase defendant107259)
(isa defendant107259 PoliceOfficer-Municipal) (isa defendant107259 SocialBeing)

;;;The defendant was in a restaurant.
(objectFoundInLocation defendant107259 restaurant107428)
(isa restaurant107428 RestaurantSpace)
      (isa  defendant107259  SocialBeing)  (relationExistsInstance  defendants
CourtCase defendant107259)

;;;The plaintiff was intoxicated.
(hasState plaintiff107470 Intoxicated)
(relationExistsInstance plaintiffs CourtCase plaintiff107470)
(isa plaintiff107470 SocialBeing)

;;;The defendant told the plaintiff the defendant was a police officer.
```

```
(recipientOfInfo tell110361 plaintiff107470)
(performedBy tell110361 defendant107259)
(isa tell110361 Informing)
(isa defendant109905 PoliceOfficer-Municipal)
(relationExistsInstance    defendants    CourtCase    defendant107259)    (isa
defendant107259 SocialBeing) (isa plaintiff107470 SocialBeing)
      (relationExistsInstance plaintiffs CourtCase plaintiff107470)

;;;Then the plaintiff pointed a gun at the defendant.
(relationExistsInstance nextInSequence List point107946)
(isa point107946 MakingAThreat)
(threateningAgent point107946 plaintiff107470)
(threatenedAgent point107946 defendant107259)
(isa gun107990 Weapon)
(doneBy point107946 plaintiff107470)
(intendedToLocation point107946 defendant107259)
(instrument-Generic point107946 gun107990)
(isa point107946 AimingSomething)
(isa plaintiff107470 SocialBeing)
(relationExistsInstance defendants CourtCase defendant107259)
(relationExistsInstance plaintiffs CourtCase plaintiff107470)
(isa defendant107259 SocialBeing) (isa gun107990 Gun)

;;;Then the plaintiff ran away.
(relationExistsInstance nextInSequence List run108489)
(isa run108489 Running) (performedBy run108489 plaintiff107470)
(isa plaintiff107470 SocialBeing)
(relationExistsInstance plaintiffs CourtCase plaintiff107470)

;;;The defendant told the plaintiff to stop running.
(isa tell108282 Requesting-CommunicationAct)
(doneBy stop108382 plaintiff107470) (isa run108489 Running)
(temporalThingTerminated stop108382 run108489)
      (performedBy run108489 plaintiff107470)
(performedBy      tell108282      defendant107259)      (isa      stop108382
DiscontinuingAnActivity)
      (situationTopic tell108282 stop108382)
        (recipientOfInfo tell108282 plaintiff107470)
      (relationExistsInstance defendants CourtCase defendant107259)
        (relationExistsInstance plaintiffs CourtCase run108489)
      (isa run108489 SocialBeing) (isa defendant107259 SocialBeing)

;;;Then the plaintiff pointed a gun at the defendant again.
(relationExistsInstance  nextInSequence  List  point108778)(isa  point108778
MakingAThreat)
(threateningAgent point108778 plaintiff107470)
(threatenedAgent point108778 defendant107259) (isa gun107990 Weapon)
(doneBy point108778 plaintiff107470)
(intendedToLocation point108778 defendant107259)
(instrument-Generic point108778 gun107990)
(isa point108778 AimingSomething)
(isa plaintiff107470 SocialBeing)
(isa gun107990 Gun)
(isa defendant107259 SocialBeing) (isa point108778 RepeatedEventType)
      (relationExistsInstance defendants CourtCase defendant107259)
        (relationExistsInstance plaintiffs CourtCase plaintiff107470)
```

```
;;;Then the defendant fired the gun at the plaintiff.
(relationExistsInstance nextInSequence List fire109059)
(isa gun107991 Weapon)
(doneBy fire109059 defendant107259)
(to-Generic fire109059 plaintiff107470)
(launcherInShooting fire109059 gun107991)
      (isa fire109059 ShootingAGun) (isa gun107991 Gun) (isa defendant107259
SocialBeing)  (isa plaintiff107470 SocialBeing)
       (relationExistsInstance plaintiffs CourtCase plaintiff107470)
      (relationExistsInstance defendants CourtCase defendant107259)
(startsAfterStartingOf fire109059 point108778)
(isa fire109059 AttackOnTangible)
(performedBy fire109059 defendant107259)
(objectAttacked fire109059 plaintiff107470)
(purposeInEvent defendant107259 fire109059 (damages fire109059
plaintiff107470))

;;;Then the plaintiff aimed the gun at the defendant again.
(relationExistsInstance nextInSequence List aim108778)
(isa aim108778 MakingAThreat) (threateningAgent aim108778 plaintiff107470)
(threatenedAgent aim108778 defendant107259) (isa gun107990 Weapon)
      (doneBy aim108778 plaintiff107470) (intendedToLocation aim108778
defendant107259) (instrument-Generic aim108778 gun107990)
      (isa aim108778 AimingSomething) (isa plaintiff107470 SocialBeing)
        (isa defendant107259 SocialBeing)
      (isa aim108778 RepeatedEventType) (isa gun107990 Gun)
        (relationExistsInstance defendants CourtCase defendant107259)
      (relationExistsInstance plaintiffs CourtCase plaintiff107470)

;;;Then the defendant shot the plaintiff again.
(relationExistsInstance nextInSequence List shoot109629)
(isa shoot109629 AttackOnTangible)
(performedBy shoot109629 defendant107259)
(objectAttacked shoot109629 plaintiff107470)
(damages shoot109629 plaintiff107470)
(intendedToLocation shoot109629 plaintiff107470)
(isa shoot109629 ShootingAndHittingSomething)
(doneBy shoot109629 defendant107259)
      (isa defendant107259 SocialBeing) (isa shoot109629 RepeatedEventType)
(relationExistsInstance plaintiffs CourtCase plaintiff107470)
      (relationExistsInstance defendants CourtCase defendant107259)
        (isa plaintiff107470 SocialBeing)

(in-microtheory (LegalCaseMtFn Banks_v_Chicago))
(genlMt (LegalCaseMtFn Banks_v_Chicago) (CaseLawCorpusMtFn Banks_v_Chicago))

(isa Banks_v_Chicago CourtCase)
(plaintiffs Banks_v_Chicago plaintiff107470)
(defendants Banks_v_Chicago defendant107259)
(isa defendant107259 Defendant)
(isa plaintiff107470 Plaintiff)

;;;The defendant was a police officer.
(isa defendant107259 PoliceOfficer-Municipal)

;;;The defendant was in a restaurant.
(objectFoundInLocation defendant107259 restaurant107428)
```

```
(isa restaurant107428 RestaurantSpace)

;;;The plaintiff was intoxicated.
(hasState plaintiff107470 Intoxicated)

;;;The defendant told the plaintiff the defendant was a police officer.
;;; need the info transferred
(recipientOfInfo tell110361 plaintiff107470)
(performedBy tell110361 defendant107259)
(isa tell110361 Informing)
(infoTransferred tell110361
      (isa defendant109905 PoliceOfficer-Municipal))

;;;Then the plaintiff pointed a gun at the defendant.
(startsAfterStartingOf point107946 tell110361)
(isa point107946 MakingAThreat)
(threateningAgent point107946 plaintiff107470)
(threatenedAgent point107946 defendant107259)
(isa gun107990 Weapon)
(doneBy point107946 plaintiff107470)
(intendedToLocation point107946 defendant107259)
(instrument-Generic point107946 gun107990)
(isa gun107990 Gun)

;;;Then the plaintiff ran away.
;;; should be Fleeing
(startsAfterStartingOf run108489 point107946)
(isa run108489 Fleeing)
(performedBy run108489 plaintiff107470)

;;;The defendant told the plaintiff to stop running.
;;; needs more structure
(implies (and (isa tell108282 Requesting-CommunicationAct)
                  (performedBy tell108282 defendant107259)
                  (situationTopic tell108282 stop108382)
                  (recipientOfInfo tell108282 plaintiff107470) )
            (and (doneBy stop108382 plaintiff107470)
                  (temporalThingTerminated stop108382 run108489)
                  (isa stop108382 DiscontinuingAnActivity)))


;;;Then the plaintiff pointed a gun at the defendant again.
(startsAfterStartingOf point108778 tell108282)
(isa point108778 MakingAThreat)
(threateningAgent point108778 plaintiff107470)
(threatenedAgent point108778 defendant107259) (isa gun107990 Weapon)
(doneBy point108778 plaintiff107470)
(intendedToLocation point108778 defendant107259)
(instrument-Generic point108778 gun107990)
(isa point108778 AimingSomething)
(isa gun107990 Gun)
(isa point108778 RepeatedEventType)

;;;Then the defendant fired the gun at the plaintiff.
(startsAfterStartingOf fire109059 point108778)
(isa fire109059 AttackOnTangible)
(performedBy fire109059 defendant107259)
```

```
(objectAttacked fire109059 plaintiff107470)
(purposeInEvent defendant107259 fire109059 (damages fire109059 plaintiff107470))
(isa gun107991 Weapon)
(doneBy fire109059 defendant107259)
(to-Generic fire109059 plaintiff107470)
(launcherInShooting fire109059 gun107991)
      (isa fire109059 ShootingAGun) (isa gun107991 Gun)

;;;Then the plaintiff aimed the gun at the defendant again.
(startsAfterStartingOf aim108778 fire109059)
(isa aim108778 MakingAThreat) (threateningAgent aim108778 plaintiff107470)
(threatenedAgent aim108778 defendant107259) (isa gun107990 Weapon)
      (doneBy aim108778 plaintiff107470) (intendedToLocation aim108778 defendant107259) (instrument-Generic aim108778 gun107990)
      (isa aim108778 AimingSomething)
      (isa aim108778 RepeatedEventType) (isa gun107990 Gun)

;;;Then the defendant shot the plaintiff again.
(startsAfterStartingOf shoot109629 aim108778)
(isa shoot109629 AttackOnTangible)
(performedBy shoot109629 defendant107259)
(objectAttacked shoot109629 plaintiff107470)
(damages shoot109629 plaintiff107470)
(intendedToLocation shoot109629 plaintiff107470)
(isa shoot109629 ShootingAndHittingSomething)
(doneBy shoot109629 defendant107259)
(isa shoot109629 RepeatedEventType)

(in-microtheory (LegalCaseConclusionMtFn Banks_v_Chicago))
(genlMt (LegalCaseConclusionMtFn Banks_v_Chicago) (LegalCaseMtFn Banks_v_Chicago))

(assaultsPartyByDoing plaintiff107470 defendant107259 aim108778)
(actingInSelfDefAgainst defendant107259 plaintiff107470 shoot109629 aim108778)

(assaultsPartyByDoing plaintiff107470 defendant107259 point108778)
(actingInSelfDefAgainst defendant107259 plaintiff107470 fire109059 point108778)

(in-microtheory (LegalCaseConclusion-NegatedMtFn Banks_v_Chicago))
(not (assaultsPartyByDoing plaintiff107470 defendant107259 aim108778))
(not (actingInSelfDefAgainst defendant107259 plaintiff107470 shoot109629 aim108778))

(not (assaultsPartyByDoing plaintiff107470 defendant107259 point108778))
(not (actingInSelfDefAgainst defendant107259 plaintiff107470 fire109059 point108778))

(in-microtheory (LegalCaseConclusion-ReversedMtFn Banks_v_Chicago))
(assaultsPartyByDoing defendant107259 plaintiff107470 aim108778)
(actingInSelfDefAgainst plaintiff107470 defendant107259 shoot109629 aim108778)

(assaultsPartyByDoing defendant107259 plaintiff107470 point108778)
(actingInSelfDefAgainst plaintiff107470 defendant107259 fire109059 point108778)
```

```
(in-microtheory (LegalCaseConclusion-NegatedReversedMtFn Banks_v_Chicago))
(not (assaultsPartyByDoing defendant107259 plaintiff107470 aim108778))
(not (actingInSelfDefAgainst plaintiff107470 defendant107259 shoot109629 aim108778))

(not (assaultsPartyByDoing defendant107259 plaintiff107470 point108778))
(not (actingInSelfDefAgainst plaintiff107470 defendant107259 fire109059 point108778))
```

(in-microtheory CaseLawCorpusMt)

(isa Benamon_v_Soo_Line_RR CaseLawCorpusCase)

(in-microtheory (CaseLawCorpusMtFn Benamon_v_Soo_Line_RR))

(caseType Benamon_v_Soo_Line_RR Trespass)

(caseValence Benamon_v_Soo_Line_RR Trespass Positive)

(caseName Benamon_v_Soo_Line_RR "Benamon v Soo Line R.R.")

(caseCourt Benamon_v_Soo_Line_RR "Appellate Court of Illinois, First District, Third Division")

(caseYear Benamon_v_Soo_Line_RR "1997")

(caseReporter Benamon_v_Soo_Line_RR "294 Ill. App. 3d 85")

(caseOriginalText Benamon_v_Soo_Line_RR "The defendants owned and maintained the railroad tracks and the right-of-way at the overpass located near 300 North Fulton in Chicago on January 13, 1992. On January 13th, 1992 at approximately 5 or 6 p.m. [the plaintiff] William was walking home after watching a basketball game at a grammar school. William was 15 years old at the time. As William passed a group of boys playing basketball in an alley, one of the boys asked him for a cigarette. William told the boy he did not smoke and continued walking. When he looked back, he saw the boys pointing at him and running towards him. William thought they might want to steal his 'Bulls' coat. William testified that he ran up the railway incline near 300 North Fulton onto the tracks by the girder to hide because he was scared and wanted to do 'anything to try to get away.' He was approximately three blocks from home but did not want the gang to see where he lived. After William thought the gang was gone, he attempted to leave the area but his foot fell between the steel grates next to one of the wooden rails. He tried but could not extricate his foot before a passing freight train partially amputated his left leg below the knee. [Testimony was controverted as to whether children regularly went on the tracks. It was incontroverted that there was graffiti on the track girders near where William was injured.]")

(caseConclusion Benamon_v_Soo_Line_RR "Plaintiff was trespassing.")

(caseConclusion Benamon_v_Soo_Line_RR "Plaintiff was not entitled to trespass due to private necessity.")

(caseConclusion Benamon_v_Soo_Line_RR "Plaintiff was not a permissive user of the track.")

(caseConclusion Benamon_v_Soo_Line_RR "Defendant was not wilfull and wanton in protecting the track.")

(caseConclusion Benamon_v_Soo_Line_RR "Defendant owed plaintiff no duty of care.")

(caseSimpleText Benamon_v_Soo_Line_RR "The plaintiff is named William. The defendant owned the railroad tracks. William was walking home. William passed a group of boys. The boys chased William. William was afraid of of the boys. William ran away from the boys. William ran onto the railroad tracks. William hid near a railroad girder. There was graffiti on the girder. After William thought the gang was gone, he tried to leave. William's foot got stuck in a grate near the railroad. William could not free his foot. A passing freight train crushed William's leg. William's leg was partially amputated below the knee.")

```
(in-microtheory (rawLanguageOutputMtFn Benamon_v_Soo_Line_RR))
(genlMt (rawLanguageOutputMtFn Benamon_v_Soo_Line_RR) (CaseLawCorpusMtFn
Benamon_v_Soo_Line_RR))

((and (by-Underspecified be3759 plaintiff3747) (nameString william6936
name3779) (performedBy name3779 plaintiff3747)
      (isa name3779 NamingSomething) (isa be3759 Situation) (isa william6936
MaleHuman) (givenNames william6936 "William")
      (relationExistsInstance plaintiffs CourtCase plaintiff3747)
(namedEntityInDiscourse william6936)
      (isa plaintiff3747 SocialBeing))
 (and (objectActedOn own3904 group-of-railroad-tracks3972)
(relationInstanceExists objectActedOn own3904 Railway)
      (relationExistsInstance defendants CourtCase defendant3893) (isa
defendant3893 SocialBeing)
      (isa group-of-railroad-tracks3972 (SetOfTypeFn Railway)) (isa own3904
OwningSomething) (doneBy own3904 defendant3893)
      (elementOf railroad-tracks3972 group-of-railroad-tracks3972) (isa
railroad-tracks3972 Railway))
 (elementOf railroad-tracks3972 group-of-railroad-tracks3972) (isa railroad-
tracks3972 Railway)
 (and (isa walk4037 Walking-Generic) (performedBy walk4037 william4007)
(eventOccursAt walk4037 (HomeFn william4007))
      (isa william4007 MaleHuman) (namedEntityInDiscourse william4007)
(givenNames william4007 "William"))
 (and (isa group4193 Group) (groupMembers group4193 boy4283) (objectPassing
pass4135 william4128)
      (fe_aggregate group4193 william4128) (isa pass4135 PassingBySomething)
(objectPassed pass4135 group4193)
      (isa william4128 MaleHuman) (givenNames william4128 "William")
(namedEntityInDiscourse william4128)
      (isa boy4283 MaleChild) (fe_person boy4283 william4128))
 (and (performedBy be3759 group-of-boy4376)
      (relationInstanceExists performedBy be3759
       (CollectionSubsetFn MaleChild
        (TheSetOf ?boy4376 (and (isa ?boy4376 MaleChild) (fe_person ?boy4376
group-of-boy4376)))))
      (isa william6936 MaleHuman) (isa group-of-boy4376 (SetOfTypeFn
MaleChild)) (keyParticipants be3759 william6936)
      (givenNames william6936 "William") (namedEntityInDiscourse william6936)
(isa be3759 PursuingSomething)
      (elementOf boy4376 group-of-boy4376) (isa boy4376 MaleChild) (fe_person
boy4376 group-of-boy4376))
 (elementOf boy4376 group-of-boy4376) (and (isa boy4376 MaleChild) (fe_person
boy4376 group-of-boy4376))
 (and (awayFromLocation run4640 group-of-boy4376)
      (relationInstanceExists awayFromLocation run4640
       (CollectionSubsetFn MaleChild (TheSetOf ?boy4700 (and (fe_person
?boy4700 william6936) (isa ?boy4700 MaleChild)))))
      (givenNames william6936 "William") (isa group-of-boy4376 (SetOfTypeFn
MaleChild)) (isa william6936 MaleHuman)
```

(namedEntityInDiscourse william6936) (performedBy run4640 william6936) (isa run4640 Running)
(elementOf boy4700 group-of-boy4376) (fe_person boy4700 william6936) (isa boy4700 MaleChild))
(elementOf boy4700 group-of-boy4376) (and (fe_person boy4700 william6936) (isa boy4700 MaleChild))
(and (toLocation run4749 group-of-railroad-tracks3972) (relationInstanceExists toLocation run4749 Railway)
(givenNames william6936 "William") (isa group-of-railroad-tracks3972 (SetOfTypeFn Railway))
(namedEntityInDiscourse william6936) (isa william6936 MaleHuman) (performedBy run4749 william6936) (isa run4749 Running) (objectFoundInLocation william6936 group-of-railroad-tracks3972)
(elementOf railroad-tracks3972 group-of-railroad-tracks3972) (isa railroad-tracks3972 Railway))
(elementOf railroad-tracks3972 group-of-railroad-tracks3972) (isa railroad-tracks3972 Railway)
(and (isa railroad4963 Railway) (nameString girder5003 railroad4963) (isa hide4851 HidingOneself)
(givenNames william6936 "William") (namedEntityInDiscourse william6936) (isa william6936 MaleHuman)
(isa girder5003 Girder))
(and (on-UnderspecifiedSurface there5053 girder5003) (isa there5053 Graffiti) (isa girder5003 Girder))
(and (transferredThing leave5633 girder5003) (attemptAtPerforming try5544 leave5633) (performedBy try5544 girder5003)
(implies (and (subjectOfMentalSituation think5274 william5261) (isa think5274 Thinking))
(and (objectMoving go5418 gang5347) (isa go5418 Translocation) (isa gang5347 Gang)))
(isa leave5633 LeavingAPlace) (isa try5544 Attempting) (after-Underspecified try5544 think5274)
(isa william5261 MaleHuman) (namedEntityInDiscourse william5261) (givenNames william5261 "William"))
(and (subjectOfMentalSituation think5274 william5261) (isa think5274 Thinking))
(and (objectMoving go5418 gang5347) (isa go5418 Translocation) (isa gang5347 Gang))
(not (and (isa william6419 MaleHuman) (givenNames william6419 "William") (namedEntityInDiscourse william6419)
(possible
(and (possessiveRelation foot7261 foot6549) (isa free6442 ReleasingFromConfinement)
(objectActedOn free6442 foot6549) (doneBy free6442 william6419)))))
(and (isa william6419 MaleHuman) (givenNames william6419 "William") (namedEntityInDiscourse william6419)
(possible
(and (possessiveRelation foot7261 foot6549) (isa free6442 ReleasingFromConfinement) (objectActedOn free6442 foot6549)
(doneBy free6442 william6419))))
(and (possessiveRelation foot7261 foot6549) (isa free6442 ReleasingFromConfinement) (objectActedOn free6442 foot6549)
(doneBy free6442 william6419))
(and (objectHarmed crush6802 leg6882) (isa leg6882 Leg) (doneBy crush6802 freight-train6782)

(isa crush6802 (CausingFn DamageOutcome)) (isa william6936 MaleHuman) (isa freight-train6782 FreightTrain)
(possessiveRelation william6936 leg6882) (namedEntityInDiscourse william6936) (givenNames william6936 "William"))
(and (isa amputate6965 (RemovalFn Limb-AnimalBodyPart)) (isa leg6942 Leg) (namedEntityInDiscourse william6936)
(isa knee7083 Knee) (under-UnderspecifiedLocation leg6942 knee7083) (isa william6936 MaleHuman)
(possessiveRelation william6936 leg6942) (givenNames william6936 "William"))
(and (isa william6936 MaleHuman) (namedEntityInDiscourse william6936) (isa group-of-boy4376 (SetOfTypeFn MaleChild))
(givenNames william6936 "William") (feelsEmotion william6936 (PositiveAmountFn Fear)) (isa boy4376 MaleChild)
(age boy4376 group-of-boy4376) (fe_person boy4376 william6936) (elementOf boy4376 group-of-boy4376)
(isa boy4376 MaleChild) (age boy4376 group-of-boy4376) (fe_person boy4376 william6936))
(elementOf boy4376 group-of-boy4376)
(and (isa boy4376 MaleChild) (age boy4376 group-of-boy4376) (fe_person boy4376 william6936))
(and (to-Generic stick7359 grate7501) (near foot7261 railroad7679) (isa stick7359 Attachment)
(isa foot7261 Foot-AnimalBodyPart) (objectOfAttachment stick7359 foot7261)
(ownerOfProprietaryThing foot7261 william6936) (isa railroad7679 Railway) (possessiveRelation william6936 foot7261)
(namedEntityInDiscourse william6936) (isa william6936 MaleHuman) (givenNames william6936 "William")
(isa grate7501 Grate-Barrier)))

;;;
(in-microtheory (cleanLanguageOutputMtFn Benamon_v_Soo_Line_RR))
(genlMt (cleanLanguageOutputMtFn Benamon_v_Soo_Line_RR) (CaseLawCorpusMtFn Benamon_v_Soo_Line_RR))

;;;The plaintiff is named William.
(by-Underspecified be3759 plaintiff3747) (nameString william6936 name3779)
(performedBy name3779 plaintiff3747)
(isa name3779 NamingSomething) (isa be3759 Situation) (isa william6936 MaleHuman) (givenNames william6936 "William")
(relationExistsInstance plaintiffs CourtCase plaintiff3747)
(namedEntityInDiscourse william6936)
(isa plaintiff3747 SocialBeing)

;;;The defendant owned the railroad tracks.
(objectActedOn own3904 group-of-railroad-tracks3972)
(relationExistsInstance defendants CourtCase defendant3893)
(isa defendant3893 SocialBeing)
(isa group-of-railroad-tracks3972 (SetOfTypeFn Railway))
(isa own3904 OwningSomething) (doneBy own3904 defendant3893)
(elementOf railroad-tracks3972 group-of-railroad-tracks3972)
(isa railroad-tracks3972 Railway)

;;;William was walking home.
(isa walk4037 Walking-Generic) (performedBy walk4037 william6936)
(intendedToLocation walk4037 (HomeFn william6936))

```
(isa william6936 MaleHuman) (namedEntityInDiscourse william6936)
(givenNames william6936 "William")

;;;William passed a group of boys.
(isa group-of-boy4376 Group) (isa group-of-boy4376 (SetOfTypeFn MaleChild))
(objectPassing pass4135 william6936)
(isa pass4135 PassingBySomething)
(objectPassed pass4135 group-of-boy4376)
(isa william6936 MaleHuman) (givenNames william6936 "William")
(namedEntityInDiscourse william6936)
(elementOf boy4376 group-of-boy4376) (isa boy4376 MaleChild)

;;;The boys chased William.
(performedBy chase4387 group-of-boy4376)
(isa william6936 MaleHuman) (isa group-of-boy4376 (SetOfTypeFn MaleChild))
(keyParticipants chase4387  william6936)
(givenNames william6936 "William") (namedEntityInDiscourse william6936)
(isa chase4387  PursuingSomething)
(elementOf boy4376 group-of-boy4376) (isa boy4376 MaleChild)
(fe_person boy4376 group-of-boy4376)

;;;William was afraid of the boys.
;;; need the "of" part
(isa william6936 MaleHuman) (namedEntityInDiscourse william6936) (isa group-
of-boy4376 (SetOfTypeFn MaleChild))
(givenNames william6936 "William") (feelsEmotion william6936 (PositiveAmountFn
Fear))
(isa boy4376 MaleChild)
(age boy4376 group-of-boy4376) (fe_person boy4376 william6936) (elementOf
boy4376 group-of-boy4376)
(isa boy4376 MaleChild) (age boy4376 group-of-boy4376) (fe_person boy4376
william6936)

;;;William ran away from the boys.
(awayFromLocation run4640 group-of-boy4376)
(givenNames william6936 "William")
(isa group-of-boy4376 (SetOfTypeFn MaleChild))
(isa william6936 MaleHuman)
(namedEntityInDiscourse william6936) (performedBy run4640 william6936) (isa
run4640 Running)
(elementOf boy4700 group-of-boy4376) (fe_person boy4700 william6936) (isa
boy4700 MaleChild)

;;;William ran onto the railroad tracks.
(toLocation run4749 group-of-railroad-tracks3972)
(givenNames william6936 "William")
(isa group-of-railroad-tracks3972 (SetOfTypeFn Railway))
(namedEntityInDiscourse william6936) (isa william6936 MaleHuman)
(performedBy run4749 william6936) (isa run4749 Running)
(elementOf railroad-tracks3972 group-of-railroad-tracks3972)
(isa railroad-tracks3972 Railway)
(objectFoundInLocation william6936 group-of-railroad-tracks3972)

;;;William hid near a railroad girder.
;;; need the near; need better rep'n of the girder.
(isa railroad-tracks3972 Railway) (isa hide4851 HidingOneself)
(givenNames william6936 "William") (namedEntityInDiscourse william6936)
```

```
(isa william6936 MaleHuman)
(isa girder5003 Girder)

;;;There was graffiti on the girder.
(on-UnderspecifiedSurface there5053 girder5003)
(isa there5053 Graffiti)
(isa girder5003 Girder)

;;;After William thought the gang was gone, he tried to leave.
;;; need the away from; need the william is the one leaving
(transferredThing leave5633 william6936)
(attemptAtPerforming try5544 leave5633) (performedBy try5544 william6936)
(implies (and (subjectOfMentalSituation think5274 william6936)
                        (isa think5274 Thinking))
            (and (objectMoving go5418 gang5347)
                  (isa go5418 Translocation)
                  (isa group-of-boy4376 Gang)))
(isa leave5633 LeavingAPlace)
(isa try5544 Attempting)
(after-Underspecified try5544 think5274)
(isa william6936 MaleHuman) (namedEntityInDiscourse william5261) (givenNames
william5261 "William")

;;;William's foot got stuck in a grate near the railroad.
(to-Generic stick7359 grate7501) (near grate7501 railroad-tracks3972)
(isa stick7359 Attachment)
(isa foot7261 Foot-AnimalBodyPart) (objectOfAttachment stick7359 foot7261)
(ownerOfProprietaryThing foot7261 william6936) (isa railroad-tracks3972
Railway) (possessiveRelation william6936 foot7261)
(namedEntityInDiscourse william6936) (isa william6936 MaleHuman)
(givenNames william6936 "William")
(isa grate7501 Grate-Barrier)

;;;William could not free his foot.
(not (and (isa william6419 MaleHuman) (givenNames william6419 "William")

                  (namedEntityInDiscourse william6419)
           (possible
            (and (possessiveRelation foot7261 foot6549)
                        (isa free6442 ReleasingFromConfinement)
                 (objectActedOn free6442 foot6549) (doneBy free6442
william6936)))))

;;;A passing freight train crushed William's leg.
(objectHarmed crush6802 leg6882) (isa leg6882 Leg) (doneBy crush6802 freight-
train6782)
      (isa crush6802 (CausingFn DamageOutcome)) (isa william6936 MaleHuman)
        (isa freight-train6782 FreightTrain)
      (possessiveRelation william6936 leg6882) (namedEntityInDiscourse
william6936) (givenNames william6936 "William")

;;;William's leg was partially amputated below the knee.
(isa amputate6965 (RemovalFn Limb-AnimalBodyPart)) (isa leg6942 Leg)
(namedEntityInDiscourse william6936)
      (isa knee7083 Knee) (under-UnderspecifiedLocation leg6942 knee7083) (isa
william6936 MaleHuman)
```

```
      (possessiveRelation   william6936   leg6942)   (givenNames   william6936
"William")

;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
(in-microtheory (cleanLanguageOutputMtFn-NoDiscourseNF Benamon_v_Soo_Line_RR))
(genlMt     (cleanLanguageOutputMtFn-NoDiscourseNF     Benamon_v_Soo_Line_RR)
(CaseLawCorpusMtFn Benamon_v_Soo_Line_RR))

;;;The plaintiff is named William.
(by-Underspecified  be3759  plaintiff3747)  (nameString  william6936  name3779)
(performedBy name3779 plaintiff3747)
(isa  name3779  NamingSomething)  (isa  be3759  Situation)  (isa  william6936
MaleHuman) (givenNames william6936 "William")
(relationExistsInstance plaintiffs CourtCase plaintiff3747)
(namedEntityInDiscourse william6936)
(isa plaintiff3747 SocialBeing)

;;;The defendant owned the railroad tracks.
(objectActedOn own3904 group-of-railroad-tracks3972)
(relationExistsInstance defendants CourtCase defendant3893)
(isa defendant3893 SocialBeing)
(isa group-of-railroad-tracks3972 (SetOfTypeFn Railway))
(isa own3904 OwningSomething) (doneBy own3904 defendant3893)
(elementOf railroad-tracks3972 group-of-railroad-tracks3972)
(isa railroad-tracks3972 Railway)

;;;William was walking home.
(isa walk4037 Walking-Generic) (performedBy walk4037 william6936)
(intendedToLocation walk4037 (HomeFn william6936))
(isa william6936 MaleHuman) (namedEntityInDiscourse william6936)
(givenNames william6936 "William")

;;;William passed a group of boys.
(isa group-of-boy4376 Group) (isa group-of-boy4376 (SetOfTypeFn MaleChild))
(objectPassing pass4135 william6936)
(isa pass4135 PassingBySomething)
(objectPassed pass4135 group-of-boy4376)
(isa william6936 MaleHuman) (givenNames william6936 "William")
(namedEntityInDiscourse william6936)
(elementOf boy4376 group-of-boy4376) (isa boy4376 MaleChild)

;;;The boys chased William.
(performedBy chase4387 group-of-boy4376)
(isa william6936 MaleHuman) (isa group-of-boy4376 (SetOfTypeFn MaleChild))
(keyParticipants chase4387  william6936)
(givenNames william6936 "William") (namedEntityInDiscourse william6936)
(isa chase4387  PursuingSomething)
(elementOf boy4376 group-of-boy4376) (isa boy4376 MaleChild)
(fe_person boy4376 group-of-boy4376)

;;;William was afraid of the boys.
;;; need the "of" part
(isa  william6936  MaleHuman)  (namedEntityInDiscourse  william6936)  (isa  group-
of-boy4376 (SetOfTypeFn MaleChild))
(givenNames william6936 "William") (feelsEmotion william6936 (PositiveAmountFn
Fear))
(isa boy4376 MaleChild)
```

```
(age boy4376 group-of-boy4376) (fe_person boy4376 william6936) (elementOf
boy4376 group-of-boy4376)
(isa boy4376 MaleChild) (age boy4376 group-of-boy4376) (fe_person boy4376
william6936)

;;;William ran away from the boys.
(awayFromLocation run4640 group-of-boy4376)
(givenNames william6936 "William")
(isa group-of-boy4376 (SetOfTypeFn MaleChild))
(isa william6936 MaleHuman)
(namedEntityInDiscourse william6936) (performedBy run4640 william6936) (isa
run4640 Running)
(elementOf boy4700 group-of-boy4376) (fe_person boy4700 william6936) (isa
boy4700 MaleChild)

;;;William ran onto the railroad tracks.
(toLocation run4749 group-of-railroad-tracks3972)
(givenNames william6936 "William")
(isa group-of-railroad-tracks3972 (SetOfTypeFn Railway))
(namedEntityInDiscourse william6936) (isa william6936 MaleHuman)
(performedBy run4749 william6936) (isa run4749 Running)
(elementOf railroad-tracks3972 group-of-railroad-tracks3972)
(isa railroad-tracks3972 Railway)
(objectFoundInLocation william6936 group-of-railroad-tracks3972)

;;;William hid near a railroad girder.
;;; need the near; need better rep'n of the girder.
(isa railroad-tracks3972 Railway) (isa hide4851 HidingOneself)
(givenNames william6936 "William") (namedEntityInDiscourse william6936)
(isa william6936 MaleHuman)
(isa girder5003 Girder)

;;;There was graffiti on the girder.
(on-UnderspecifiedSurface there5053 girder5003)
(isa there5053 Graffiti)
(isa girder5003 Girder)

;;;After William thought the gang was gone, he tried to leave.
;;; need the away from; need the william is the one leaving
(transferredThing leave5633 william6936)
(attemptAtPerforming try5544 leave5633) (performedBy try5544 william6936)
(implies (and (subjectOfMentalSituation think5274 william6936)
                        (isa think5274 Thinking))
            (and (objectMoving go5418 gang5347)
                  (isa go5418 Translocation)
                  (isa group-of-boy4376 Gang)))
(isa leave5633 LeavingAPlace)
(isa try5544 Attempting)
(after-Underspecified try5544 think5274)
(isa william6936 MaleHuman) (namedEntityInDiscourse william5261) (givenNames
william5261 "William")

;;;William's foot got stuck in a grate near the railroad.
(to-Generic stick7359 grate7501) (near grate7501 railroad-tracks3972)
(isa stick7359 Attachment)
(isa foot7261 Foot-AnimalBodyPart) (objectOfAttachment stick7359 foot7261)
```

```
(ownerOfProprietaryThing   foot7261   william6936)   (isa   railroad-tracks3972
Railway) (possessiveRelation william6936 foot7261)
(namedEntityInDiscourse william6936) (isa william6936 MaleHuman)
(givenNames william6936 "William")
(isa grate7501 Grate-Barrier)

;;;William could not free his foot.
(not (and (isa william6419 MaleHuman) (givenNames william6419 "William")

                  (namedEntityInDiscourse william6419)
           (possible
            (and (possessiveRelation foot7261 foot6549)
                       (isa free6442 ReleasingFromConfinement)
                 (objectActedOn    free6442    foot6549)    (doneBy    free6442
william6936)))))

;;;A passing freight train crushed William's leg.
(objectHarmed crush6802 leg6882) (isa leg6882 Leg) (doneBy crush6802 freight-
train6782)
      (isa crush6802 (CausingFn DamageOutcome)) (isa william6936 MaleHuman)
        (isa freight-train6782 FreightTrain)
      (possessiveRelation    william6936    leg6882)    (namedEntityInDiscourse
william6936) (givenNames william6936 "William")

;;;William's leg was partially amputated below the knee.
(isa   amputate6965   (RemovalFn   Limb-AnimalBodyPart))   (isa   leg6942   Leg)
(namedEntityInDiscourse william6936)
      (isa knee7083 Knee) (under-UnderspecifiedLocation leg6942 knee7083) (isa
william6936 MaleHuman)
      (possessiveRelation   william6936   leg6942)   (givenNames   william6936
"William")

;;;;;;;;;;;;;;;;;;;;;;

(in-microtheory (LegalCaseMtFn Benamon_v_Soo_Line_RR))
(genlMt     (LegalCaseMtFn     Benamon_v_Soo_Line_RR)     (CaseLawCorpusMtFn
Benamon_v_Soo_Line_RR))

(isa Benamon_v_Soo_Line_RR CourtCase)
(plaintiffs Benamon_v_Soo_Line_RR william6936)
(defendants Benamon_v_Soo_Line_RR defendant3893)
(isa defendant3893 Defendant)
(isa william6936 Plaintiff)

;;;The plaintiff is named William.
(nameString william6936 "William")
(isa william6936 MaleHuman) (givenNames william6936 "William")
(relationExistsInstance plaintiffs CourtCase william6936)
(namedEntityInDiscourse william6936)

;;;The defendant owned the railroad tracks.
(objectActedOn own3904 group-of-railroad-tracks3972)
(isa group-of-railroad-tracks3972 (SetOfTypeFn Railway))
(isa own3904 OwningSomething) (doneBy own3904 defendant3893)
(elementOf railroad-tracks3972 group-of-railroad-tracks3972)
(isa railroad-tracks3972 Railway)
```

```
;;;William was walking home.
(isa walk4037 Walking-Generic) (performedBy walk4037 william6936)
(intendedToLocation walk4037 (HomeFn william6936))

;;;William passed a group of boys.
(isa group-of-boy4376 Group) (isa group-of-boy4376 (SetOfTypeFn MaleChild))
(objectPassing pass4135 william6936)
(isa pass4135 PassingBySomething)
(objectPassed pass4135 group-of-boy4376)
(elementOf boy4376 group-of-boy4376) (isa boy4376 MaleChild)

;;;The boys chased William.
(performedBy chase4387 group-of-boy4376)
(keyParticipants chase4387  william6936)
(isa chase4387  PursuingSomething)

;;;William was afraid of the boys.
;;; adding the hasEmotionAbout
(hasEmotionAbout william6936 group-of-boy4376)
(feelsEmotion william6936 (PositiveAmountFn Fear))

;;;William ran away from the boys.
(awayFromLocation run4640 group-of-boy4376)
(performedBy run4640 william6936) (isa run4640 Running)

;;;William ran onto the railroad tracks.
;;; adding a temporal ordering relation
(startsAfterStartingOf run4749 chase4387)
(toLocation run4749 group-of-railroad-tracks3972)
(performedBy run4749 william6936)
(isa run4749 Running)
(objectFoundInLocation william6936 group-of-railroad-tracks3972)

;;;William hid near a railroad girder.
;;; need role relations on hide; need better rep'n of the girder.
(isa hide4851 HidingOneself)
(doneBy hide4851 william6936)
(near william6936 girder5003)
(isa girder5003 Girder)

;;;There was graffiti on the girder.
(on-UnderspecifiedSurface graffiti5053 girder5003)
(isa graffiti5053 Graffiti)

;;;After William thought the gang was gone, he tried to leave.
;;; need the away from; need the william is the one leaving
(transferredThing leave5633 william6936)
(attemptAtPerforming try5544 leave5633) (performedBy try5544 william6936)
(implies (and (subjectOfMentalSituation think5274 william6936)
                        (isa think5274 Thinking))
            (and (objectMoving go5418 gang5347)
                  (isa go5418 Translocation)
                  (isa group-of-boy4376 Gang)
                  (awayFromLocation go5418 group-of-railroad-tracks3972)))
(isa leave5633 LeavingAPlace)
(isa try5544 Attempting)
(after-Underspecified try5544 think5274)
```

```
;;;William's foot got stuck in a grate near the railroad.
(to-Generic stick7359 grate7501) (near grate7501 railroad-tracks3972)
(isa stick7359 Attachment)
(isa foot7261 Foot-AnimalBodyPart) (objectOfAttachment stick7359 foot7261)
(ownerOfProprietaryThing foot7261 william6936)
(possessiveRelation william6936 foot7261)
(isa grate7501 Grate-Barrier)

;;;William could not free his foot.
(not (possible (and (possessiveRelation foot7261 foot6549)
                              (isa free6442 ReleasingFromConfinement)
                              (objectActedOn free6442 foot6549)
                              (doneBy free6442 william6936))))

;;;A passing freight train crushed William's leg.
(objectHarmed crush6802 leg6882) (isa leg6882 Leg)
(doneBy crush6802 freight-train6782)
(isa crush6802 (CausingFn DamageOutcome))
(isa freight-train6782 FreightTrain)
(possessiveRelation william6936 leg6882)

;;;William's leg was partially amputated below the knee.
;;; need it to be William's leg being amputated; william's knee
(isa amputate6965 (RemovalFn Limb-AnimalBodyPart))
(objectActedOn amputate6965 leg6942)
(isa leg6942 Leg)
(isa knee7083 Knee) (under-UnderspecifiedLocation amputate6965 knee7083)
(possessiveRelation william6936 leg6942)
(possessiveRelation william6936 knee7083)

(in-microtheory (LegalCaseConclusionMtFn Benamon_v_Soo_Line_RR))
(genlMt    (LegalCaseConclusionMtFn    Benamon_v_Soo_Line_RR)    (LegalCaseMtFn
Benamon_v_Soo_Line_RR))

(trespassOnPropertyByAction william6936 railroad-tracks3972 run4749)
(trespassOnPropertyByAction william6936 railroad-tracks3972 hide4851)

(not (trespassJustifiedByNecessity william6936 run4749 chase4387))
(not (permissiveOrLicensedUserOfProperty william6936 railroad-tracks3972))
(not (willfullWantonActivity defendant3893 crush6802))
(not (owedDutyOfCare defendant3893 william6936))

(in-microtheory (LegalCaseConclusion-NegatedMtFn Benamon_v_Soo_Line_RR))
(not (trespassOnPropertyByAction william6936 railroad-tracks3972 run4749))
(not (trespassOnPropertyByAction william6936 railroad-tracks3972 hide4851))

(trespassJustifiedByNecessity william6936 run4749 chase4387)
(permissiveOrLicensedUserOfProperty william6936 railroad-tracks3972)
(willfullWantonActivity defendant3893 crush6802)
(owedDutyOfCare defendant3893 william6936)

(in-microtheory (LegalCaseConclusion-ReversedMtFn Benamon_v_Soo_Line_RR))
(trespassOnPropertyByAction defendant3893 railroad-tracks3972 run4749)
(trespassOnPropertyByAction defendant3893 railroad-tracks3972 hide4851)

(not (trespassJustifiedByNecessity defendant3893 run4749 chase4387))
```

```
(not (permissiveOrLicensedUserOfProperty defendant3893 railroad-tracks3972))
(not (willfullWantonActivity william6936 crush6802))
(not (owedDutyOfCare william6936 defendant3893))

(in-microtheory                    (LegalCaseConclusion-NegatedReversedMtFn
Benamon_v_Soo_Line_RR))
(not (trespassOnPropertyByAction defendant3893 railroad-tracks3972 run4749))
(not (trespassOnPropertyByAction defendant3893 railroad-tracks3972 hide4851))

(trespassJustifiedByNecessity defendant3893 run4749 chase4387)
(permissiveOrLicensedUserOfProperty defendant3893 railroad-tracks3972)
(willfullWantonActivity william6936 crush6802)
(owedDutyOfCare william6936 defendant3893)
```

APPENDIX B: DETAILED EXPERIMENTAL RESULTS – ALL COMPARISONS PERFORMED FOR EXPERIMENTS IN CHAPTER 6

Techniques were compared using proportion tests. The first two columns show the methods being compared; the following two columns show the number of cases each method got correct (always out of 100 cases); the third column shows the p-statistic; the fourth column shows the z-score; and the fifth column shows the p-value for the difference. The final column is provided to let the reader rapidly hone in on which results were significant.

Method 1	Method 2	True1	True2	p	z	p-value	Sig?
RRLG	GPT-J-PT@1	47	12	0.295	5.426846	2.87E-08	TRUE
RRLG	GPT-J-PT@6	47	61	0.54	-1.98627	0.023502	TRUE
RRLG	GPT-J-ST@1	47	12	0.295	5.426846	2.87E-08	TRUE
RRLG	GPT-J-ST@6	47	52	0.495	-0.70714	0.239739	f
RRLG	legalBERT	47	33	0.4	2.020726	0.021654	TRUE
RRLG	ARPG PT 2CIs @1	47	41	0.44	0.854704	0.196357	f
RRLG	ARPG PT 2CIs @6	47	71	0.59	-3.45047	0.00028	TRUE
RRLG	ARPG PT 1CIs @1	47	36	0.415	1.578613	0.057212	f
RRLG	ARPG PT 1CIs @6	47	58	0.525	-1.55758	0.059666	f
RRLG	ARPG PT 0CIs @1	47	32	0.395	2.169702	0.015015	TRUE
RRLG	ARPG PT 0CIs @6	47	43	0.45	0.568535	0.284836	f
RRLG	ARPG ST 2CIs @1	47	32	0.395	2.169702	0.015015	TRUE
RRLG	ARPG ST 2CIs @6	47	58	0.525	-1.55758	0.059666	f
RRLG	ARPG ST 1CIs @1	47	29	0.38	2.622224	0.004368	TRUE
RRLG	ARPG ST 1CIs @6	47	47	0.47	0	0.5	f
RRLG	ARPG ST 0CIs @1	47	28	0.375	2.775128	0.002759	TRUE
RRLG	ARPG ST 0CIs @6	47	35	0.41	1.725235	0.042243	TRUE
RRLG	PAPR PT @1	47	28	0.375	2.775128	0.002759	TRUE
RRLG	PAPR PT @6	47	72	0.595	-3.60113	0.000158	TRUE
RRLG	PAPR ST @1	47	17	0.32	4.547543	2.71E-06	TRUE
RRLG	PAPR ST @6	47	46	0.465	0.141769	0.443631	f
PAPR ST @6	GPT-J-PT@1	46	12	0.29	5.29829	5.84E-08	TRUE
PAPR ST @6	GPT-J-PT@6	46	61	0.535	-2.12654	0.016729	TRUE
PAPR ST @6	GPT-J-ST@1	46	12	0.29	5.29829	5.84E-08	TRUE
PAPR ST @6	GPT-J-ST@6	46	52	0.49	-0.8487	0.198025	f
PAPR ST @6	legalBERT	46	33	0.395	1.880408	0.030026	TRUE

PAPR ST @6	ARPG PT 2CIs @1	46	41	0.435	0.713159	0.237874	f
PAPR ST @6	ARPG PT 2CIs @6	46	71	0.585	-3.58776	0.000167	TRUE
PAPR ST @6	ARPG PT 1CIs @1	46	36	0.41	1.437696	0.07526	f
PAPR ST @6	ARPG PT 1CIs @6	46	58	0.52	-1.69842	0.044715	TRUE
PAPR ST @6	ARPG PT 0CIs @1	46	32	0.39	2.029625	0.021197	TRUE
PAPR ST @6	ARPG PT 0CIs @6	46	43	0.445	0.426854	0.334743	f
PAPR ST @6	ARPG ST 2CIs @1	46	32	0.39	2.029625	0.021197	TRUE
PAPR ST @6	ARPG ST 2CIs @6	46	58	0.52	-1.69842	0.044715	TRUE
PAPR ST @6	ARPG ST 1CIs @1	46	29	0.375	2.483009	0.006514	TRUE
PAPR ST @6	ARPG ST 1CIs @6	46	47	0.465	-0.14177	0.443631	f
PAPR ST @6	ARPG ST 0CIs @1	46	28	0.37	2.636249	0.004191	TRUE
PAPR ST @6	ARPG ST 0CIs @6	46	35	0.405	1.584498	0.05654	f
PAPR ST @6	PAPR PT @1	46	28	0.37	2.636249	0.004191	TRUE
PAPR ST @6	PAPR PT @6	46	72	0.59	-3.73801	9.27E-05	TRUE
PAPR ST @6	PAPR ST @1	46	17	0.315	4.414509	5.06E-06	TRUE
PAPR ST @1	GPT-J-PT@1	17	12	0.145	1.004125	0.157659	f
PAPR ST @1	GPT-J-PT@6	17	61	0.39	-6.37882	8.92E-11	TRUE
PAPR ST @1	GPT-J-ST@1	17	12	0.145	1.004125	0.157659	f
PAPR ST @1	GPT-J-ST@6	17	52	0.345	-5.20622	9.64E-08	TRUE
PAPR ST @1	legalBERT	17	33	0.25	-2.61279	0.00449	TRUE
PAPR ST @1	ARPG PT 2CIs @1	17	41	0.29	-3.73997	9.2E-05	TRUE
PAPR ST @1	ARPG PT 2CIs @6	17	71	0.44	-7.69234	7.22E-15	TRUE
PAPR ST @1	ARPG PT 1CIs @1	17	36	0.265	-3.04419	0.001167	TRUE
PAPR ST @1	ARPG PT 1CIs @6	17	58	0.375	-5.98843	1.06E-09	TRUE
PAPR ST @1	ARPG PT 0CIs @1	17	32	0.245	-2.46615	0.006829	TRUE
PAPR ST @1	ARPG PT 0CIs @6	17	43	0.3	-4.01189	3.01E-05	TRUE
PAPR ST @1	ARPG ST 2CIs @1	17	32	0.245	-2.46615	0.006829	TRUE
PAPR ST @1	ARPG ST 2CIs @6	17	58	0.375	-5.98843	1.06E-09	TRUE
PAPR ST @1	ARPG ST 1CIs @1	17	29	0.23	-2.01631	0.021884	TRUE
PAPR ST @1	ARPG ST 1CIs @6	17	47	0.32	-4.54754	2.71E-06	TRUE
PAPR ST @1	ARPG ST 0CIs @1	17	28	0.225	-1.86267	0.031254	TRUE
PAPR ST @1	ARPG ST 0CIs @6	17	35	0.26	-2.90172	0.001856	TRUE
PAPR ST @1	PAPR PT @1	17	28	0.225	-1.86267	0.031254	TRUE
PAPR ST @1	PAPR PT @6	17	72	0.445	-7.82566	2.55E-15	TRUE
PAPR PT @6	GPT-J-PT@1	72	12	0.42	8.596024	0	TRUE
PAPR PT @6	GPT-J-PT@6	72	61	0.665	1.647952	0.049681	TRUE
PAPR PT @6	GPT-J-ST@1	72	12	0.42	8.596024	0	TRUE
PAPR PT @6	GPT-J-ST@6	72	52	0.62	2.913583	0.001787	TRUE
PAPR PT @6	legalBERT	72	33	0.525	5.52234	1.67E-08	TRUE
PAPR PT @6	ARPG PT 2CIs @1	72	41	0.565	4.421584	4.9E-06	TRUE
PAPR PT @6	ARPG PT 2CIs @6	72	71	0.715	0.156642	0.437763	f
PAPR PT @6	ARPG PT 1CIs @1	72	36	0.54	5.107539	1.63E-07	TRUE
PAPR PT @6	ARPG PT 1CIs @6	72	58	0.65	2.075498	0.01897	TRUE

PAPR PT @6	ARPG PT 0CIs @1	72	32	0.52	5.661385	7.51E-09	TRUE
PAPR PT @6	ARPG PT 0CIs @6	72	43	0.575	4.148152	1.68E-05	TRUE
PAPR PT @6	ARPG ST 2CIs @1	72	32	0.52	5.661385	7.51E-09	TRUE
PAPR PT @6	ARPG ST 2CIs @6	72	58	0.65	2.075498	0.01897	TRUE
PAPR PT @6	ARPG ST 1CIs @1	72	29	0.505	6.081422	5.96E-10	TRUE
PAPR PT @6	ARPG ST 1CIs @6	72	47	0.595	3.601132	0.000158	TRUE
PAPR PT @6	ARPG ST 0CIs @1	72	28	0.5	6.22254	2.45E-10	TRUE
PAPR PT @6	ARPG ST 0CIs @6	72	35	0.535	5.245457	7.79E-08	TRUE
PAPR PT @6	PAPR PT @1	72	28	0.5	6.22254	2.45E-10	TRUE
PAPR PT @1	GPT-J-PT@1	28	12	0.2	2.828427	0.002339	TRUE
PAPR PT @1	GPT-J-PT@6	28	61	0.445	-4.6954	1.33E-06	TRUE
PAPR PT @1	GPT-J-ST@1	28	12	0.2	2.828427	0.002339	TRUE
PAPR PT @1	GPT-J-ST@6	28	52	0.4	-3.4641	0.000266	TRUE
PAPR PT @1	legalBERT	28	33	0.305	-0.76791	0.221269	f
PAPR PT @1	ARPG PT 2CIs @1	28	41	0.345	-1.93374	0.026573	TRUE
PAPR PT @1	ARPG PT 2CIs @6	28	71	0.495	-6.08142	5.96E-10	TRUE
PAPR PT @1	ARPG PT 1CIs @1	28	36	0.32	-1.21268	0.112626	f
PAPR PT @1	ARPG PT 1CIs @6	28	58	0.43	-4.28484	9.14E-06	TRUE
PAPR PT @1	ARPG PT 0CIs @1	28	32	0.3	-0.61721	0.268547	f
PAPR PT @1	ARPG PT 0CIs @6	28	43	0.355	-2.21657	0.013326	TRUE
PAPR PT @1	ARPG ST 2CIs @1	28	32	0.3	-0.61721	0.268547	f
PAPR PT @1	ARPG ST 2CIs @6	28	58	0.43	-4.28484	9.14E-06	TRUE
PAPR PT @1	ARPG ST 1CIs @1	28	29	0.285	-0.15664	0.437763	f
PAPR PT @1	ARPG ST 1CIs @6	28	47	0.375	-2.77513	0.002759	TRUE
PAPR PT @1	ARPG ST 0CIs @1	28	28	0.28	0	0.5	f
PAPR PT @1	ARPG ST 0CIs @6	28	35	0.315	-1.06557	0.143309	f
ARPG ST 0CIs @6	GPT-J-PT@1	35	12	0.235	3.835731	6.26E-05	TRUE
ARPG ST 0CIs @6	GPT-J-PT@6	35	61	0.48	-3.6799	0.000117	TRUE
ARPG ST 0CIs @6	GPT-J-ST@1	35	12	0.235	3.835731	6.26E-05	TRUE
ARPG ST 0CIs @6	GPT-J-ST@6	35	52	0.435	-2.42474	0.00766	TRUE
ARPG ST 0CIs @6	legalBERT	35	33	0.34	0.298541	0.382645	f
ARPG ST 0CIs @6	ARPG PT 2CIs @1	35	41	0.38	-0.87407	0.191039	f
ARPG ST 0CIs @6	ARPG PT 2CIs @6	35	71	0.53	-5.10036	1.7E-07	TRUE
ARPG ST 0CIs @6	ARPG PT 1CIs @1	35	36	0.355	-0.14777	0.441262	f
ARPG ST 0CIs @6	ARPG PT 1CIs @6	35	58	0.465	-3.26069	0.000556	TRUE
ARPG ST 0CIs @6	ARPG PT 0CIs @1	35	32	0.335	0.449441	0.326557	f
ARPG ST 0CIs @6	ARPG PT 0CIs @6	35	43	0.39	-1.15979	0.123068	f
ARPG ST 0CIs @6	ARPG ST 2CIs @1	35	32	0.335	0.449441	0.326557	f
ARPG ST 0CIs @6	ARPG ST 2CIs @6	35	58	0.465	-3.26069	0.000556	TRUE
ARPG ST 0CIs @6	ARPG ST 1CIs @1	35	29	0.32	0.909509	0.181541	f
ARPG ST 0CIs @6	ARPG ST 1CIs @6	35	47	0.41	-1.72524	0.042243	TRUE
ARPG ST 0CIs @6	ARPG ST 0CIs @1	35	28	0.315	1.065571	0.143309	f
ARPG ST 0CIs @1	GPT-J-PT@1	28	12	0.2	2.828427	0.002339	TRUE

ARPG ST 0CIs @1	GPT-J-PT@6	28	61	0.445	-4.6954	1.33E-06	TRUE
ARPG ST 0CIs @1	GPT-J-ST@1	28	12	0.2	2.828427	0.002339	TRUE
ARPG ST 0CIs @1	GPT-J-ST@6	28	52	0.4	-3.4641	0.000266	TRUE
ARPG ST 0CIs @1	legalBERT	28	33	0.305	-0.76791	0.221269	f
ARPG ST 0CIs @1	ARPG PT 2CIs @1	28	41	0.345	-1.93374	0.026573	TRUE
ARPG ST 0CIs @1	ARPG PT 2CIs @6	28	71	0.495	-6.08142	5.96E-10	TRUE
ARPG ST 0CIs @1	ARPG PT 1CIs @1	28	36	0.32	-1.21268	0.112626	f
ARPG ST 0CIs @1	ARPG PT 1CIs @6	28	58	0.43	-4.28484	9.14E-06	TRUE
ARPG ST 0CIs @1	ARPG PT 0CIs @1	28	32	0.3	-0.61721	0.268547	f
ARPG ST 0CIs @1	ARPG PT 0CIs @6	28	43	0.355	-2.21657	0.013326	TRUE
ARPG ST 0CIs @1	ARPG ST 2CIs @1	28	32	0.3	-0.61721	0.268547	f
ARPG ST 0CIs @1	ARPG ST 2CIs @6	28	58	0.43	-4.28484	9.14E-06	TRUE
ARPG ST 0CIs @1	ARPG ST 1CIs @1	28	29	0.285	-0.15664	0.437763	f
ARPG ST 0CIs @1	ARPG ST 1CIs @6	28	47	0.375	-2.77513	0.002759	TRUE
ARPG ST 1CIs @6	GPT-J-PT@1	47	12	0.295	5.426846	2.87E-08	TRUE
ARPG ST 1CIs @6	GPT-J-PT@6	47	61	0.54	-1.98627	0.023502	TRUE
ARPG ST 1CIs @6	GPT-J-ST@1	47	12	0.295	5.426846	2.87E-08	TRUE
ARPG ST 1CIs @6	GPT-J-ST@6	47	52	0.495	-0.70714	0.239739	f
ARPG ST 1CIs @6	legalBERT	47	33	0.4	2.020726	0.021654	TRUE
ARPG ST 1CIs @6	ARPG PT 2CIs @1	47	41	0.44	0.854704	0.196357	f
ARPG ST 1CIs @6	ARPG PT 2CIs @6	47	71	0.59	-3.45047	0.00028	TRUE
ARPG ST 1CIs @6	ARPG PT 1CIs @1	47	36	0.415	1.578613	0.057212	f
ARPG ST 1CIs @6	ARPG PT 1CIs @6	47	58	0.525	-1.55758	0.059666	f
ARPG ST 1CIs @6	ARPG PT 0CIs @1	47	32	0.395	2.169702	0.015015	TRUE
ARPG ST 1CIs @6	ARPG PT 0CIs @6	47	43	0.45	0.568535	0.284836	f
ARPG ST 1CIs @6	ARPG ST 2CIs @1	47	32	0.395	2.169702	0.015015	TRUE
ARPG ST 1CIs @6	ARPG ST 2CIs @6	47	58	0.525	-1.55758	0.059666	f
ARPG ST 1CIs @6	ARPG ST 1CIs @1	47	29	0.38	2.622224	0.004368	TRUE
ARPG ST 1CIs @1	GPT-J-PT@1	29	12	0.205	2.977648	0.001452	TRUE
ARPG ST 1CIs @1	GPT-J-PT@6	29	61	0.45	-4.54828	2.7E-06	TRUE
ARPG ST 1CIs @1	GPT-J-ST@1	29	12	0.205	2.977648	0.001452	TRUE
ARPG ST 1CIs @1	GPT-J-ST@6	29	52	0.405	-3.31304	0.000461	TRUE
ARPG ST 1CIs @1	legalBERT	29	33	0.31	-0.61156	0.270414	f
ARPG ST 1CIs @1	ARPG PT 2CIs @1	29	41	0.35	-1.779	0.03762	TRUE
ARPG ST 1CIs @1	ARPG PT 2CIs @6	29	71	0.5	-5.9397	1.43E-09	TRUE
ARPG ST 1CIs @1	ARPG PT 1CIs @1	29	36	0.325	-1.05679	0.145303	f
ARPG ST 1CIs @1	ARPG PT 1CIs @6	29	58	0.435	-4.13632	1.76E-05	TRUE
ARPG ST 1CIs @1	ARPG PT 0CIs @1	29	32	0.305	-0.46075	0.32249	f
ARPG ST 1CIs @1	ARPG PT 0CIs @6	29	43	0.36	-2.06239	0.019585	TRUE
ARPG ST 1CIs @1	ARPG ST 2CIs @1	29	32	0.305	-0.46075	0.32249	f
ARPG ST 1CIs @1	ARPG ST 2CIs @6	29	58	0.435	-4.13632	1.76E-05	TRUE
ARPG ST 2CIs @6	GPT-J-PT@1	58	12	0.35	6.819494	4.57E-12	TRUE
ARPG ST 2CIs @6	GPT-J-PT@6	58	61	0.595	-0.43214	0.332821	f

ARPG ST 2CIs @6	GPT-J-ST@1	58	12	0.35	6.819494	4.57E-12	TRUE
ARPG ST 2CIs @6	GPT-J-ST@6	58	52	0.55	0.852803	0.196884	f
ARPG ST 2CIs @6	legalBERT	58	33	0.455	3.54994	0.000193	TRUE
ARPG ST 2CIs @6	ARPG PT 2CIs @1	58	41	0.495	2.404283	0.008102	TRUE
ARPG ST 2CIs @6	ARPG PT 2CIs @6	58	71	0.645	-1.92103	0.027364	TRUE
ARPG ST 2CIs @6	ARPG PT 1CIs @1	58	36	0.47	3.116885	0.000914	TRUE
ARPG ST 2CIs @6	ARPG PT 1CIs @6	58	58	0.58	0	0.5	f
ARPG ST 2CIs @6	ARPG PT 0CIs @1	58	32	0.45	3.695479	0.00011	TRUE
ARPG ST 2CIs @6	ARPG PT 0CIs @6	58	43	0.505	2.121426	0.016943	TRUE
ARPG ST 2CIs @6	ARPG ST 2CIs @1	58	32	0.45	3.695479	0.00011	TRUE
ARPG ST 2CIs @1	GPT-J-PT@1	32	12	0.22	3.413944	0.00032	TRUE
ARPG ST 2CIs @1	GPT-J-PT@6	32	61	0.465	-4.1113	1.97E-05	TRUE
ARPG ST 2CIs @1	GPT-J-ST@1	32	12	0.22	3.413944	0.00032	TRUE
ARPG ST 2CIs @1	GPT-J-ST@6	32	52	0.42	-2.86534	0.002083	TRUE
ARPG ST 2CIs @1	legalBERT	32	33	0.325	-0.15097	0.44	f
ARPG ST 2CIs @1	ARPG PT 2CIs @1	32	41	0.365	-1.32189	0.093103	f
ARPG ST 2CIs @1	ARPG PT 2CIs @6	32	71	0.515	-5.51792	1.72E-08	TRUE
ARPG ST 2CIs @1	ARPG PT 1CIs @1	32	36	0.34	-0.59708	0.275227	f
ARPG ST 2CIs @1	ARPG PT 1CIs @6	32	58	0.45	-3.69548	0.00011	TRUE
ARPG ST 2CIs @1	ARPG PT 0CIs @1	32	32	0.32	0	0.5	f
ARPG ST 2CIs @1	ARPG PT 0CIs @6	32	43	0.375	-1.60665	0.054065	f
ARPG PT 0CIs @6	GPT-J-PT@1	43	12	0.275	4.909207	4.57E-07	TRUE
ARPG PT 0CIs @6	GPT-J-PT@6	43	61	0.52	-2.54762	0.005423	TRUE
ARPG PT 0CIs @6	GPT-J-ST@1	43	12	0.275	4.909207	4.57E-07	TRUE
ARPG PT 0CIs @6	GPT-J-ST@6	43	52	0.475	-1.27439	0.101263	f
ARPG PT 0CIs @6	legalBERT	43	33	0.38	1.456791	0.072587	f
ARPG PT 0CIs @6	ARPG PT 2CIs @1	43	41	0.42	0.286534	0.387235	f
ARPG PT 0CIs @6	ARPG PT 2CIs @6	43	71	0.57	-3.99918	3.18E-05	TRUE
ARPG PT 0CIs @6	ARPG PT 1CIs @1	43	36	0.395	1.012527	0.155643	f
ARPG PT 0CIs @6	ARPG PT 1CIs @6	43	58	0.505	-2.12143	0.016943	TRUE
ARPG PT 0CIs @6	ARPG PT 0CIs @1	43	32	0.375	1.606653	0.054065	f
ARPG PT 0CIs @1	GPT-J-PT@1	32	12	0.22	3.413944	0.00032	TRUE
ARPG PT 0CIs @1	GPT-J-PT@6	32	61	0.465	-4.1113	1.97E-05	TRUE
ARPG PT 0CIs @1	GPT-J-ST@1	32	12	0.22	3.413944	0.00032	TRUE
ARPG PT 0CIs @1	GPT-J-ST@6	32	52	0.42	-2.86534	0.002083	TRUE
ARPG PT 0CIs @1	legalBERT	32	33	0.325	-0.15097	0.44	f
ARPG PT 0CIs @1	ARPG PT 2CIs @1	32	41	0.365	-1.32189	0.093103	f
ARPG PT 0CIs @1	ARPG PT 2CIs @6	32	71	0.515	-5.51792	1.72E-08	TRUE
ARPG PT 0CIs @1	ARPG PT 1CIs @1	32	36	0.34	-0.59708	0.275227	f
ARPG PT 0CIs @1	ARPG PT 1CIs @6	32	58	0.45	-3.69548	0.00011	TRUE
ARPG PT 1CIs @6	GPT-J-PT@1	58	12	0.35	6.819494	4.57E-12	TRUE
ARPG PT 1CIs @6	GPT-J-PT@6	58	61	0.595	-0.43214	0.332821	f
ARPG PT 1CIs @6	GPT-J-ST@1	58	12	0.35	6.819494	4.57E-12	TRUE

ARPG PT 1CIs @6	GPT-J-ST@6	58	52	0.55	0.852803	0.196884	f
ARPG PT 1CIs @6	legalBERT	58	33	0.455	3.54994	0.000193	TRUE
ARPG PT 1CIs @6	ARPG PT 2CIs @1	58	41	0.495	2.404283	0.008102	TRUE
ARPG PT 1CIs @6	ARPG PT 2CIs @6	58	71	0.645	-1.92103	0.027364	TRUE
ARPG PT 1CIs @6	ARPG PT 1CIs @1	58	36	0.47	3.116885	0.000914	TRUE
ARPG PT 1CIs @1	GPT-J-PT@1	36	12	0.24	3.973597	3.54E-05	TRUE
ARPG PT 1CIs @1	GPT-J-PT@6	36	61	0.485	-3.53713	0.000202	TRUE
ARPG PT 1CIs @1	GPT-J-ST@1	36	12	0.24	3.973597	3.54E-05	TRUE
ARPG PT 1CIs @1	GPT-J-ST@6	36	52	0.44	-2.27921	0.011327	TRUE
ARPG PT 1CIs @1	legalBERT	36	33	0.345	0.446248	0.327709	f
ARPG PT 1CIs @1	ARPG PT 2CIs @1	36	41	0.385	-0.72659	0.23374	f
ARPG PT 1CIs @1	ARPG PT 2CIs @6	36	71	0.535	-4.96192	3.49E-07	TRUE
ARPG PT 2CIs @6	GPT-J-PT@1	71	12	0.415	8.467107	0	TRUE
ARPG PT 2CIs @6	GPT-J-PT@6	71	61	0.66	1.492704	0.067757	f
ARPG PT 2CIs @6	GPT-J-ST@1	71	12	0.415	8.467107	0	TRUE
ARPG PT 2CIs @6	GPT-J-ST@6	71	52	0.615	2.761027	0.002881	TRUE
ARPG PT 2CIs @6	legalBERT	71	33	0.52	5.378316	3.76E-08	TRUE
ARPG PT 2CIs @6	ARPG PT 2CIs @1	71	41	0.56	4.273522	9.62E-06	TRUE
ARPG PT 2CIs @1	GPT-J-PT@1	41	12	0.265	4.646398	1.69E-06	TRUE
ARPG PT 2CIs @1	GPT-J-PT@6	41	61	0.51	-2.82899	0.002335	TRUE
ARPG PT 2CIs @1	GPT-J-ST@1	41	12	0.265	4.646398	1.69E-06	TRUE
ARPG PT 2CIs @1	GPT-J-ST@6	41	52	0.465	-1.55946	0.059444	f
ARPG PT 2CIs @1	legalBERT	41	33	0.37	1.171666	0.120666	f
legalBERT	GPT-J-PT@1	33	12	0.225	3.556004	0.000188	TRUE
legalBERT	GPT-J-PT@6	33	61	0.47	-3.96694	3.64E-05	TRUE
legalBERT	GPT-J-ST@1	33	12	0.225	3.556004	0.000188	TRUE
legalBERT	GPT-J-ST@6	33	52	0.425	-2.71775	0.003286	TRUE
GPT-J-ST@6	GPT-J-PT@1	52	12	0.32	6.063391	6.66E-10	TRUE
GPT-J-ST@6	GPT-J-PT@6	52	61	0.565	-1.28369	0.099626	f
GPT-J-ST@6	GPT-J-ST@1	52	12	0.32	6.063391	6.66E-10	TRUE
GPT-J-ST@1	GPT-J-PT@1	12	12	0.12	0	0.5	f
GPT-J-ST@1	GPT-J-PT@6	12	61	0.365	-7.19694	3.08E-13	TRUE
GPT-J-PT@6	GPT-J-PT@1	61	12	0.365	7.196938	3.08E-13	TRUE

Precision@6 testing. These analyses feature more granular comparisons by doctrine, of positive vs. negative performance, and more.

Cases	Method 1	Method 2	n1	T1	n2	T2	p	z	p-value	Sig?
All	ARPG_0CIs	ARPG_1CIs	100	35	100	47	0.41	-1.72524	0.042243	TRUE
Assault	ARPG_0CIs	ARPG_1CIs	17	6	17	9	0.441176	-1.03619	0.150057	f
Battery	ARPG_0CIs	ARPG_1CIs	40	10	40	16	0.325	-1.43223	0.076039	f

Trespass	ARPG_0CIs	ARPG_1CIs	43	19	43	22	0.476744	-0.6477	0.25859	f
All	ARPG_1CIs	ARPG_2CIs	100	47	100	58	0.525	-1.55758	0.059666	f
Assault	ARPG_1CIs	ARPG_2CIs	17	9	17	9	0.529412	0	0.5	f
Battery	ARPG_1CIs	ARPG_2CIs	40	16	40	22	0.475	-1.34332	0.089584	f
Trespass	ARPG_1CIs	ARPG_2CIs	43	22	43	27	0.569767	-1.08898	0.138081	f
All	ARPG_0CIs	ARPG_2CIs	100	35	100	58	0.465	-3.26069	0.000556	TRUE
Assault	ARPG_0CIs	ARPG_2CIs	17	6	17	9	0.441176	-1.03619	0.150057	f
Battery	ARPG_0CIs	ARPG_2CIs	40	10	40	22	0.4	-2.73861	0.003085	TRUE
Trespass	ARPG_0CIs	ARPG_2CIs	43	19	43	27	0.534884	-1.72954	0.041856	TRUE
All	ARPG_0CIs_PT	ARPG_1CIs_PT	100	43	100	58	0.505	-2.12143	0.016943	TRUE
Assault	ARPG_0CIs_PT	ARPG_1CIs_PT	17	7	17	10	0.5	-1.02899	0.151742	f
Battery	ARPG_0CIs_PT	ARPG_1CIs_PT	40	12	40	20	0.4	-1.82574	0.033945	TRUE
Trespass	ARPG_0CIs_PT	ARPG_1CIs_PT	43	24	43	28	0.604651	-0.8822	0.188834	f
All	ARPG_1CIs_PT	ARPG_2CIs_PT	100	58	100	71	0.645	-1.92103	0.027364	TRUE
Assault	ARPG_1CIs_PT	ARPG_2CIs_PT	17	10	17	10	0.588235	0	0.5	f
Battery	ARPG_1CIs_PT	ARPG_2CIs_PT	40	20	40	26	0.575	-1.35699	0.087392	f
Trespass	ARPG_1CIs_PT	ARPG_2CIs_PT	43	28	43	35	0.732558	-1.70535	0.044065	TRUE
All	ARPG_0CIs_PT	ARPG_2CIs_PT	100	43	100	71	0.57	-3.99918	3.18E-05	TRUE
Assault	ARPG_0CIs_PT	ARPG_2CIs_PT	17	7	17	10	0.5	-1.02899	0.151742	f
Battery	ARPG_0CIs_PT	ARPG_2CIs_PT	40	12	40	26	0.475	-3.13442	0.000861	TRUE
Trespass	ARPG_0CIs_PT	ARPG_2CIs_PT	43	24	43	35	0.686047	-2.55584	0.005297	TRUE
All	ARPG_0CIs	PAPR	100	35	100	46	0.405	-1.5845	0.05654	f
Assault	ARPG_0CIs	PAPR	17	6	17	11	0.5	-1.71499	0.043174	TRUE
Battery	ARPG_0CIs	PAPR	40	10	40	20	0.375	-2.3094	0.010461	TRUE
Trespass	ARPG_0CIs	PAPR	43	19	43	15	0.395349	0.882202	0.188834	f
All	ARPG_0CIs_PT	PAPR_PT	100	43	100	72	0.575	-4.14815	1.68E-05	TRUE
Assault	ARPG_0CIs_PT	PAPR_PT	17	7	17	16	0.676471	-3.2993	0.000485	TRUE
Battery	ARPG_0CIs_PT	PAPR_PT	40	12	40	27	0.4875	-3.35515	0.000397	TRUE
Trespass	ARPG_0CIs_PT	PAPR_PT	43	24	43	29	0.616279	-1.10873	0.133774	f
All	ARPG_1CIs	PAPR	100	47	100	46	0.465	0.141769	0.443631	f
Assault	ARPG_1CIs	PAPR	17	9	17	11	0.588235	-0.69693	0.242923	f
Battery	ARPG_1CIs	PAPR	40	16	40	20	0.45	-0.89893	0.184344	f
Trespass	ARPG_1CIs	PAPR	43	22	43	15	0.430233	1.524573	0.063683	f
All	ARPG_2CIs	PAPR	100	58	100	46	0.52	1.698416	0.044715	TRUE
Assault	ARPG_2CIs	PAPR	17	9	17	11	0.588235	-0.69693	0.242923	f
Battery	ARPG_2CIs	PAPR	40	22	40	20	0.525	0.447774	0.327158	f
Trespass	ARPG_2CIs	PAPR	43	27	43	15	0.488372	2.588687	0.004817	TRUE
All	ARPG_2CIs_PT	PAPR_PT	100	71	100	72	0.715	-0.15664	0.437763	f
Assault	ARPG_2CIs_PT	PAPR_PT	17	10	17	16	0.764706	-2.42582	0.007637	TRUE
Battery	ARPG_2CIs_PT	PAPR_PT	40	26	40	27	0.6625	-0.23644	0.406545	f
Trespass	ARPG_2CIs_PT	PAPR_PT	43	35	43	29	0.744186	1.482857	0.069056	f
All	ARPG_0CIs	GPT_J_ST	100	35	100	52	0.435	-2.42474	0.00766	TRUE
Assault	ARPG_0CIs	GPT_J_ST	17	6	17	6	0.352941	0	0.5	f
Battery	ARPG_0CIs	GPT_J_ST	40	10	40	20	0.375	-2.3094	0.010461	TRUE
Trespass	ARPG_0CIs	GPT_J_ST	43	19	43	26	0.523256	-1.51129	0.065357	f
All	ARPG_0CIs_PT	GPT_J_PT	100	43	100	61	0.52	-2.54762	0.005423	TRUE

Assault	ARPG_0CIs_PT	GPT_J_PT	17	7	17	6	0.382353	0.352905	0.36208	f
Battery	ARPG_0CIs_PT	GPT_J_PT	40	12	40	23	0.4375	-2.47912	0.006585	TRUE
Trespass	ARPG_0CIs_PT	GPT_J_PT	43	24	43	32	0.651163	-1.81002	0.035146	TRUE
All	ARPG_2CIs	GPT_J_ST	100	58	100	52	0.55	0.852803	0.196884	f
Assault	ARPG_2CIs	GPT_J_ST	17	9	17	6	0.441176	1.036187	0.150057	f
Battery	ARPG_2CIs	GPT_J_ST	40	22	40	20	0.525	0.447774	0.327158	f
Trespass	ARPG_2CIs	GPT_J_ST	43	27	43	26	0.616279	0.221745	0.412256	f
All	ARPG_2CIs_PT	GPT_J_PT	100	71	100	61	0.66	1.492704	0.067757	f
Assault	ARPG_2CIs_PT	GPT_J_PT	17	10	17	6	0.470588	1.374369	0.084664	f
Battery	ARPG_2CIs_PT	GPT_J_PT	40	26	40	23	0.6125	0.688474	0.245577	f
Trespass	ARPG_2CIs_PT	GPT_J_PT	43	35	43	32	0.77907	0.779752	0.217769	f
All	ARPG_0CIs	legalBERT	100	35	99	33	0.341709	0.247858	0.402122	f
Assault	ARPG_0CIs	legalBERT	17	6	17	9	0.441176	-1.03619	0.150057	f
Battery	ARPG_0CIs	legalBERT	40	10	39	14	0.303797	-1.05296	0.146179	f
Trespass	ARPG_0CIs	legalBERT	43	19	43	10	0.337209	2.052841	0.020044	TRUE
All	ARPG_0CIs_PT	legalBERT	100	43	99	33	0.38191	1.403336	0.080258	f
Assault	ARPG_0CIs_PT	legalBERT	17	7	17	9	0.470588	-0.68718	0.245983	f
Battery	ARPG_0CIs_PT	legalBERT	40	12	39	14	0.329114	-0.55772	0.288519	f
Trespass	ARPG_0CIs_PT	legalBERT	43	24	43	10	0.395349	3.087707	0.001009	TRUE
All	ARPG_2CIs	legalBERT	100	58	99	33	0.457286	3.492384	0.000239	TRUE
Assault	ARPG_2CIs	legalBERT	17	9	17	9	0.529412	0	0.5	f
Battery	ARPG_2CIs	legalBERT	40	22	39	14	0.455696	1.704441	0.044149	TRUE
Trespass	ARPG_2CIs	legalBERT	43	27	43	10	0.430233	3.702536	0.000107	TRUE
All	ARPG_2CIs_PT	legalBERT	100	71	99	33	0.522613	5.318913	5.22E-08	TRUE
Assault	ARPG_2CIs_PT	legalBERT	17	10	17	9	0.558824	0.345396	0.364898	f
Battery	ARPG_2CIs_PT	legalBERT	40	26	39	14	0.506329	2.586692	0.004845	TRUE
Trespass	ARPG_2CIs_PT	legalBERT	43	35	43	10	0.523256	5.39748	3.38E-08	TRUE
All	PAPR	GPT_J_ST	100	46	100	52	0.49	-0.8487	0.198025	f
Assault	PAPR	GPT_J_ST	17	11	17	6	0.5	1.714986	0.043174	TRUE
Battery	PAPR	GPT_J_ST	40	20	40	20	0.5	0	0.5	f
Trespass	PAPR	GPT_J_ST	43	15	43	26	0.476744	-2.37489	0.008777	TRUE
All	PAPR_PT	GPT_J_PT	100	72	100	61	0.665	1.647952	0.049681	TRUE
Assault	PAPR_PT	GPT_J_PT	17	16	17	6	0.647059	3.588703	0.000166	TRUE
Battery	PAPR_PT	GPT_J_PT	40	27	40	23	0.625	0.92376	0.177806	f
Trespass	PAPR_PT	GPT_J_PT	43	29	43	32	0.709302	-0.71242	0.238103	f
All	PAPR	legalBERT	100	46	99	33	0.396985	1.826006	0.033925	TRUE
Assault	PAPR	legalBERT	17	11	17	9	0.588235	0.696932	0.242923	f
Battery	PAPR	legalBERT	40	20	39	14	0.43038	1.265693	0.102812	f
Trespass	PAPR	legalBERT	43	15	43	19	0.395349	-0.8822	0.188834	f
All	PAPR_PT	legalBERT	100	72	99	33	0.527638	5.462888	2.34E-08	TRUE
Assault	PAPR_PT	legalBERT	17	16	17	9	0.735294	2.721111	0.003253	TRUE
Battery	PAPR_PT	legalBERT	40	27	39	14	0.518987	2.8107	0.002472	TRUE
Trespass	PAPR_PT	legalBERT	43	29	43	19	0.55814	2.171385	0.014951	TRUE
Pos	ARPG_0CIs	GPT_J_ST	71	7	71	44	0.359155	-6.47202	4.83E-11	TRUE
Neg	ARPG_0CIs	GPT_J_ST	29	28	29	8	0.62069	5.412294	3.11E-08	TRUE
Pos	ARPG_0CIs_PT	GPT_J_PT	71	15	71	52	0.471831	-6.21982	2.49E-10	TRUE

Neg	ARPG_0CIs_PT	GPT_J_PT	29	28	29	9	0.637931	5.191072	1.05E-07	TRUE
Pos	ARPG_2CIs	GPT_J_ST	71	30	71	44	0.521127	-2.35181	0.009341	TRUE
Neg	ARPG_2CIs	GPT_J_ST	29	28	29	8	0.62069	5.412294	3.11E-08	TRUE
Pos	ARPG_2CIs_PT	GPT_J_PT	71	43	71	52	0.669014	-1.605	0.054247	f
Neg	ARPG_2CIs_PT	GPT_J_PT	29	28	29	9	0.637931	5.191072	1.05E-07	TRUE
Pos	ARPG_0CIs	legalBERT	71	7	70	23	0.212766	-3.33623	0.000425	TRUE
Neg	ARPG_0CIs	legalBERT	29	28	29	10	0.655172	4.972556	3.3E-07	TRUE
Pos	ARPG_0CIs_PT	legalBERT	71	15	70	23	0.269504	-1.5696	0.058254	f
Neg	ARPG_0CIs_PT	legalBERT	29	28	29	10	0.655172	4.972556	3.3E-07	TRUE
Pos	ARPG_2CIs	legalBERT	71	30	70	23	0.375887	1.151779	0.124706	f
Neg	ARPG_2CIs	legalBERT	29	28	29	10	0.655172	4.972556	3.3E-07	TRUE
Pos	ARPG_2CIs_PT	legalBERT	71	43	70	23	0.468085	3.296573	0.000489	TRUE
Neg	ARPG_2CIs_PT	legalBERT	29	28	29	10	0.655172	4.972556	3.3E-07	TRUE
Pos	PAPR	GPT_J_ST	71	38	71	44	0.577465	-1.01933	0.154024	f
Neg	PAPR	GPT_J_ST	29	8	29	8	0.275862	0	0.5	f
Pos	PAPR_PT	GPT_J_PT	71	53	71	52	0.739437	0.191183	0.424191	f
Neg	PAPR_PT	GPT_J_PT	29	19	29	9	0.482759	2.627691	0.004298	TRUE
Pos	PAPR	legalBERT	71	38	70	23	0.432624	2.476235	0.006639	TRUE
Neg	PAPR	legalBERT	29	8	29	10	0.310345	-0.56765	0.285138	f
Pos	PAPR_PT	legalBERT	71	53	70	23	0.539007	4.977421	3.22E-07	TRUE
Neg	PAPR_PT	legalBERT	29	19	29	10	0.5	2.363516	0.009051	TRUE
	Pos	Neg	71	7	29	28	0.35	-8.24744	0	TRUE
	Pos	Neg	71	15	29	28	0.43	-6.91308	2.37E-12	TRUE
	Pos	Neg	71	30	29	28	0.58	-4.992	2.99E-07	TRUE
	Pos	Neg	71	43	29	28	0.71	-3.59883	0.00016	TRUE
	Pos	Neg	71	38	29	8	0.46	2.361225	0.009107	TRUE
	Pos	Neg	71	53	29	19	0.72	0.922751	0.178069	f
	Pos	Neg	71	44	29	8	0.52	3.123078	0.000895	TRUE
	Pos	Neg	71	52	29	9	0.61	3.926399	4.31E-05	TRUE
	Pos	Neg	70	23	29	10	0.333333	-0.15615	0.437956	f
All	RRLG	ARPG0CIs	100	47	100	35	0.41	1.725235	0.042243	TRUE
Assault	RRLG	ARPG0CIs	17	8	17	6	0.411765	0.696932	0.242923	f
Battery	RRLG	ARPG0CIs	40	21	40	10	0.3875	2.524404	0.005795	TRUE
Trespass	RRLG	ARPG0CIs	43	18	43	19	0.430233	-0.2178	0.413794	f
Pos	RRLG	ARPG0CIs	71	23	71	7	0.211268	3.289232	0.000502	TRUE
Neg	RRLG	ARPG0CIs	29	24	29	28	0.896552	-1.72463	0.042297	TRUE
All	RRLG	ARPG2CIs_PT	100	47	100	71	0.59	-3.45047	0.00028	TRUE
Assault	RRLG	ARPG2CIs_PT	17	8	17	10	0.529412	-0.68718	0.245983	f
Battery	RRLG	ARPG2CIs_PT	40	21	40	26	0.5875	-1.13556	0.128071	f
Trespass	RRLG	ARPG2CIs_PT	43	18	43	35	0.616279	-3.76967	8.17E-05	TRUE
Pos	RRLG	ARPG2CIs_PT	71	23	71	43	0.464789	-3.36508	0.000383	TRUE
Neg	RRLG	ARPG2CIs_PT	29	24	29	28	0.896552	-1.72463	0.042297	TRUE
All	RRLG	PAPR	100	47	100	46	0.465	0.141769	0.443631	f
Assault	RRLG	PAPR	17	8	17	11	0.558824	-1.03619	0.150057	f
Battery	RRLG	PAPR	40	21	40	20	0.5125	0.223677	0.411504	f
Trespass	RRLG	PAPR	43	18	43	15	0.383721	0.665236	0.25295	f

Pos	RRLG	PAPR	71	23	71	38	0.429577	-2.54289	0.005497	TRUE
Neg	RRLG	PAPR	29	24	29	8	0.551724	4.224471	1.2E-05	TRUE
All	RRLG	PAPR_PT	100	47	100	72	0.595	-3.60113	0.000158	TRUE
Assault	RRLG	PAPR_PT	17	8	17	16	0.705882	-3.01109	0.001302	TRUE
Battery	RRLG	PAPR_PT	40	21	40	27	0.6	-1.36931	0.085452	f
Trespass	RRLG	PAPR_PT	43	18	43	29	0.546512	-2.38265	0.008594	TRUE
Pos	RRLG	PAPR_PT	71	23	71	53	0.535211	-5.04762	2.24E-07	TRUE
Neg	RRLG	PAPR_PT	29	24	29	19	0.741379	1.499354	0.066891	f
All	RRLG	legalBERT	100	47	99	33	0.40201	1.966022	0.024648	TRUE
Assault	RRLG	legalBERT	17	8	17	9	0.5	-0.343	0.3658	f
Battery	RRLG	legalBERT	40	21	39	14	0.443038	1.48522	0.068743	f
Trespass	RRLG	legalBERT	43	18	43	10	0.325581	1.840968	0.032813	TRUE
Pos	RRLG	legalBERT	71	23	70	23	0.326241	-0.0586	0.476634	f
Neg	RRLG	legalBERT	29	24	29	10	0.586207	3.732475	9.48E-05	TRUE
All	RRLG	GPT_J_Strict_Truth	100	47	100	52	0.495	-0.70714	0.239739	f
Assault	RRLG	GPT_J_Strict_Truth	17	8	17	6	0.411765	0.696932	0.242923	f
Battery	RRLG	GPT_J_Strict_Truth	40	21	40	20	0.5125	0.223677	0.411504	f
Trespass	RRLG	GPT_J_Strict_Truth	43	18	43	26	0.511628	-1.72579	0.042192	TRUE
Pos	RRLG	GPT_J_Strict_Truth	71	23	71	44	0.471831	-3.53017	0.000208	TRUE
Neg	RRLG	GPT_J_Strict_Truth	29	24	29	8	0.551724	4.224471	1.2E-05	TRUE
RRLG	RRLG_Pos	RRLG_Neg	71	23	29	24	0.47	-4.57893	2.34E-06	TRUE
All	RRLG	GPT_J_PT_Truth	100	47	100	61	0.54	-1.98627	0.023502	TRUE
Assault	RRLG	GPT_J_PT_Truth	17	8	17	6	0.411765	0.696932	0.242923	f
Battery	RRLG	GPT_J_PT_Truth	40	21	40	23	0.55	-0.44947	0.326548	f
Trespass	RRLG	GPT_J_PT_Truth	43	18	43	32	0.581395	-3.06014	0.001106	TRUE
Pos	RRLG	GPT_J_PT_Truth	71	23	71	52	0.528169	-4.87499	5.44E-07	TRUE
Neg	RRLG	GPT_J_PT_Truth	29	24	29	9	0.568966	3.977208	3.49E-05	TRUE
All	PAPR	PAPR_PT	100	46	100	72	0.59	-3.73801	9.27E-05	TRUE
Assault	PAPR	PAPR_PT	17	11	17	16	0.794118	-2.1207	0.016974	TRUE
Battery	PAPR	PAPR_PT	40	20	40	27	0.5875	-1.58978	0.055942	f
Trespass	PAPR	PAPR_PT	43	15	43	29	0.511628	-3.02013	0.001263	TRUE
Pos	PAPR	PAPR_PT	71	38	71	53	0.640845	-2.62379	0.004348	TRUE
Neg	PAPR	PAPR_PT	29	8	29	19	0.465517	-2.89564	0.001892	TRUE
All	ARPG2CIs	ARPG2CIs_PT	100	58	100	71	0.645	-1.92103	0.027364	TRUE

Assault	ARPG2CIs	ARPG2CIs_PT	17	9	17	10	0.558824	-0.3454	0.364898	f
Battery	ARPG2CIs	ARPG2CIs_PT	40	22	40	26	0.6	-0.91287	0.180655	f
Trespass	ARPG2CIs	ARPG2CIs_PT	43	27	43	35	0.72093	-1.92326	0.027224	TRUE
All	ARPG1CIs	ARPG1CIs_PT	100	47	100	58	0.525	-1.55758	0.059666	f
Assault	ARPG1CIs	ARPG1CIs_PT	17	9	17	10	0.558824	-0.3454	0.364898	f
Battery	ARPG1CIs	ARPG1CIs_PT	40	16	40	20	0.45	-0.89893	0.184344	f
Trespass	ARPG1CIs	ARPG1CIs_PT	43	22	43	28	0.581395	-1.31149	0.094847	f
All	ARPG0CIs	ARPG0CIs_PT	100	35	100	43	0.39	-1.15979	0.123068	f
Assault	ARPG0CIs	ARPG0CIs_PT	17	6	17	7	0.382353	-0.35291	0.36208	f
Battery	ARPG0CIs	ARPG0CIs_PT	40	10	40	12	0.275	-0.50078	0.308262	f
Trespass	ARPG0CIs	ARPG0CIs_PT	43	19	43	24	0.5	-1.07833	0.140444	f
All	ARPG_2CIs	PAPR_PT	100	58	100	72	0.65	-2.0755	0.01897	TRUE
Assault	ARPG_2CIs	PAPR_PT	17	9	17	16	0.735294	-2.72111	0.003253	TRUE
Battery	ARPG_2CIs	PAPR_PT	40	22	40	27	0.6125	-1.14746	0.125597	f
Trespass	ARPG_2CIs	PAPR_PT	43	27	43	29	0.651163	-0.45251	0.325452	f
All	ARPG1CIs_PT	PAPR_PT	100	58	100	72	0.65	-2.0755	0.01897	TRUE
Assault	ARPG1CIs_PT	PAPR_PT	17	10	17	16	0.764706	-2.42582	0.007637	TRUE
Battery	ARPG1CIs_PT	PAPR_PT	40	20	40	27	0.5875	-1.58978	0.055942	f
Trespass	ARPG1CIs_PT	PAPR_PT	43	28	43	29	0.662791	-0.22809	0.409787	f