Predicting Trade Secret Case Outcomes using Argument Schemes and Learned Quantitative Value Effect Tradeoffs

Matthias Grabmair
Language Technologies Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA, USA

ABSTRACT

This paper presents the Value Judgment Formalism and its experimental implementation in the VJAP system, which is capable of arguing about, and predicting outcomes of, a set of trade secret misappropriation cases. VJAP creates an argument graph for each case using argument schemes and a representation of values underlying trade secret law and effects of facts on these values. It balances effects on values in each case and analogizes it to tradeoffs in precedents. It predicts case outcomes using a confidence measure computed from the graph and generates textual legal arguments justifying its predictions. The confidence propagation uses quantitative weights learned from past cases using an iterative optimization method. Prediction performance on a limited dataset is competitive with common machine learning models. The results and VJAP's behavior are discussed in detail.

CCS CONCEPTS

• **Applied computing → Law**; Decision analysis;

KEYWORDS

Artificial Intelligence & law, Case-based Reasoning, Legal Reasoning, Computational Models of Argument, Machine Learning

ACM Reference format:
Matthias Grabmair. 2017. Predicting Trade Secret Case Outcomes using Argument Schemes and Learned Quantitative Value Effect Tradeoffs. In *Proceedings of ICAIL '17, London, United Kingdom, June 12-16, 2017,* 10 pages.
https://doi.org/10.1145/3086512.3086521

1 INTRODUCTION

1.1 General

One leading vision of AI&Law is to assist lawyers in researching, drafting and evaluating arguments in a professional setting. In their landmark paper, Berman & Hafner [10] elaborated on the need for case-based legal reasoning systems to incorporate a representation of the purposes of the law in order to be able to reason about their domain on a deeper level. Various approaches have been made to explore this problem through descriptive formalisms (e.g. [7]) of purposive legal reasoning, computational models of argument that take into account values (e.g. [8]), as well as experimental intelligent legal reasoning and argumentation systems (e.g. [13]).

This paper contributes to this discussion by reporting on an experiment in implementing and evaluating our previously published Value Judgment Formalism [19] (VJF) in a computer program that innovates by using said formalism to generate arguments and predictions for a set of trade secret misappropriation cases.[1] We call the system VJAP for *Value Judgment-based Argumentative Prediction.* It creates an argument graph for each case using argument schemes and a representation of values underlying trade secret law as well as effects of facts on these values. It balances effects on values in each case and analogizes it to tradeoffs in precedents. It predicts case outcomes using a confidence measure computed from the graph and generates textual legal arguments justifying its predictions. The confidence propagation uses quantitative weights learned from past cases using an iterative optimization method. Prediction performance on a limited dataset is competitive with common machine learning models.

The remainder of the introduction explains the VJF's legal theory assumptions. Section 2 outlines the formalism. Section 3 discusses its computational implementation into the VJAP system. Section 4 presents the evaluation experiment followed by an in-depth discussion. Section 6, 7 and 8 elaborate on related work followed by future work and conclusory remarks.

1.2 Legal Theory Assumptions

The VJAP experiment starts from a set of legal theory assumptions in order to develop a formalism of purposive case-based legal argumentation and implement it computationally. The resulting system then can be evaluated in terms of its capacity to generate intelligent legal arguments and predict case outcomes.

Legal Argumentation is about Balancing Values: Complex regulatory systems may lead to arguably inequitable results if applied too literally. To remedy this, legal argumentation provides tools with which an equitable outcome in such "hard cases" [16] can be justified while staying inside the boundaries provided by statutes, codes, regulations, and precedent cases.

Positive Law Guides Balancing: Balancing arguments must be properly integrated with a legal system's mechanics to facilitate equitable outcomes in cases. *Legal argumentation* means advocating that a certain solution/decision to a given problem (i.e. sustaining some interests at the expense of others) is more or less in accordance with the coordination of interests contained in the applicable legal norms. The VJF refers to these decisions as *value judgments* or

ICAIL '17, June 12-16, 2017, London, United Kingdom

[1] This paper is a condensed version of [17].

tradeoffs. All nontrivial legal reasoning (i.e. everything beyond straightforward backward chaining of rules) is an interplay between legal norms and their underlying interest coordination.

Unified Terminology: Individual and collective interests underlying legal reasoning are typically termed "values", "principles", "interests", "purposes", and the like. To develop a formal model we use the term "values" as a label for the corresponding elements which stand in for the interests of the concerned legal subjects without subdividing them into a more fine-grained typology. Recent work [2] has related our conception of values and value judgments to Eisenberg's 'social propositions'.

2 THE VALUE JUDGMENT FORMALISM

2.1 The Core Formalism

The VJF assumes a model of defeasible reasoning with argument schemes. An argument is a statement that a set of premise propositions warrants a conclusion. An argument scheme is an argument blueprint that can be instantiated. A premise of a scheme can either be established as a fact or taken from the conclusion of another argument, thereby gradually constructing an argument graph structure. We denote an argument as $a_1 \wedge \ldots \wedge a_n \rightarrow c$, where $a_1, \ldots, a_n$ is a set of premise propositions and c is the argument's conclusion proposition. An argument for c is an argument against $\neg c$ and vice versa. $\neg p$ represents the non-monotonic negation of p.

Definition 2.1. A **fact pattern** is an atomic or compound proposition in a formal language representing a part of the domain of discourse. Let F be the set of possible fact patterns closed under union. In an adversarial setting of a plaintiff π and defendant δ, facts that favor one side or another are denoted f^π or f^δ if needed.

Definition 2.2. A **rule** is a proposition of the form $f_1 \Rightarrow f_2$, where $f_1, f_2 \in F$, assigns a conclusion to an antecedent.

Definition 2.3. A **value** is a legal concept abstracting a set of one or more interests of an individual agent or groups of agents in the legal system such that one can speak of a change in a certain fact pattern as promoting or demoting the given value. Let V be the set of values in the domain of discourse.

Definition 2.4. A **situation** is a tuple $\langle F_s, R \rangle$ where $F_s \subseteq F$ are the given facts and R are applicable rules. Let S be the set of possible situations. If $s \in S$ and proposition x is either a fact or rule, then we denote the new situation as $s' = s \cup x$.

Definition 2.5. If an argument for fact f can be constructed using the available schemes and knowledge in a situation s or not, $s \vdash f$ and $s \nvdash f$ denote situation s as **argumentatively entailing** fact f, or not. An argument for $s \vdash f$ is an argument against $s \nvdash f$ and vice versa.

2.2 Quantitative Value Effects and Tradeoffs

Qualitative value effects as in the VJF's original model [19] are difficult to operationalize because they complicate resolving conflicting arguments and are difficult to learn from case data. We hence supplement the formalism with quantitative value effects.

Definition 2.6 (Adversarial Value Effects). Assume $v \subseteq V$ and $s \in S$. Plaintiff and defendant are arguing for an outcome fact o^π, o^δ, respectively. For plaintiff-favoring facts the function

$$\theta^+ : v, s, f^\pi, o^\pi \rightarrow [0, 1]$$

represents the **positive fact effect** of f^π in s on v if outcome o^π were decided by the judge. Analogously, the output of the function

$$\theta^- : v, s, f^\pi, o^\delta \rightarrow [-1, 0]$$

represents the **negative fact effect** of f^π in s on v if outcome o^δ were decided by the judge. These functions are complementary, i.e. $\theta^+(v, s, f^\pi, o^\pi) = |\theta^-(v, s, f^\pi, o^\delta)|$. We denote the set of all positive fact effects of deciding o^π for a given set of facts F^π across all values as $E^+_{o^\pi}(V, S, F^\pi)$ and the set of all negative effect weights of deciding o^δ as $E^-_{o^\delta}(V, S, F^\pi)$. Corresponding defendant versions of these definitions are presumed to be apparent.

A value judgment expresses that the positive effects of adding some fact to a situation outweigh its negative effects, or vice versa.[2] When deciding the main outcome of a case, the *scope* of the value judgment is global. We assume, however, that legal argumentation lets value-based concerns flow into the case assessment while not deviating too far from the guidelines provided by positive law. The existence of a rule governing the global case outcome creates a new value judgment whose scope is an antecedent of the rule.

The following definitions assume a plaintiff and defendant party.

Definition 2.7. Assume situation $s \in S$, $F^\pi, F^\delta \subset F$ and applicable values V. A **value judgment** (or **tradeoff**) in favor of the plaintiff is a proposition comparing value effect sets of the form

$$E^+_{o^\pi}(V, s, F^\pi) > E^-_{o^\pi}(V, s, F^\delta)$$

If adequate, this can be shortened to $E^+_{o^\pi}(s, F^\pi) > E^-_{o^\pi}(s, F^\delta)$. An inverted version exists for the defendant. A **scoped** value judgment in favor of the plaintiff is a value judgment

$$E^+_{c^\pi}(V, s, F^\pi) > E^-_{c^\pi}(V, s, F^\delta)$$

where c is not the global decision outcome but an intermediate legal concept, which is an antecedent of some rule $r \in s$ which is necessary for a global outcome in favor of the plaintiff. A symmetric scheme exists for the defendant.

This scoping effect of value judgments cascades down the tree of auxiliary rules and their antecedents. Skillful legal argumentation includes the ability to spot the 'weakest link' in this chain with the best chances of persuading the decision body to agree with a certain locally scoped value judgment whose effect propagates up to the global assessment of the case.

3 IMPLEMENTING THE FORMALISM

VJAP implements the VJF in the domain of trade secret law. It uses the dataset developed over the course of work on HYPO [4], CATO [3] and IBP [5]. In the dataset, cases are represented as sets of *factors* from a total of 26 possible factors (i.e. factual propositions in the VJF) that correspond to relevant patterns of facts in the trade secret domain that tend to favor one side or the other. There are 13 factors each favoring plaintiff and favoring defendant.

[2] This is conceptually similar to Alexy's "Weight Formula" [1], which is an arithmetic model of balancing competing legal values in a given case context as a function of the abstract weight of a principle and the degree of interference in the case.

3.1 The Trade Secret Domain

In the trade secret cases of our dataset, plaintiff and defendant are competing producers of goods. Plaintiff has developed certain product information which the defendant obtains in some way and uses to develop and market a competing product. The plaintiff then sues the defendant for misappropriation of the trade secret by an alleged illegitimate use of the product information. Substantive American trade secret law is a combination of statutory law and case law. The latter has been consolidated in the Restatement (First) of Torts §757, which many courts have adopted. The Restatement does not state a clear rule about what constitutes a trade secret, but offers more information in a commentary (emphasis supplied).

> Restatement of Torts §757: "One who discloses or uses another's trade secret, without a privilege to do so, is liable to the other if (a) he discovered the secret by improper means, or (b) his disclosure or use constitutes a breach of confidence reposed in him by the other in disclosing the secret to him, ..."
>
> Comment b. Definition of trade secret:
> "A trade secret may consist of any formula, pattern, device or compilation of **information which is used in one's business, and which gives him an opportunity to obtain an advantage over competitors** who do not know or use it. [...] An exact definition of a trade secret is not possible. Some factors to be considered in determining whether given information is one's trade secret are:
> (1) the extent to which the information is known outside of his business;
> (2) the extent to which it is known by employees and others involved in his business;
> (3) the extent of of measures taken by him to guard the secrecy of the information;
> (4) the value of the information to him and to his competitors;
> (5) the amount of effort or money expended by him in developing the information;
> (6) the ease or difficulty with which the information could be properly acquired or duplicated by others."

It follows from the emphasized passages that, first, the information needs to be valuable by conveying a competitive advantage and, second, the inventor must make efforts to keep it secret. VJAP models these Restatement and commentary rules in its domain model (see Fig. 1). They form a tree structure of framed *issue propositions* connected by logical and/or. The factors that make up the individual cases are associated with the 'leaf' issues, which are special because they are reasoned about by taking recourse to precedent cases. Other issues are the main issue (*trade-secret-misappropriation-claim*) or intermediate issues (*info-trade-secret* , *improper-means* , *info-misappropriated*) and established from Restatement rules. All issues are anchors for scoped value judgments.

Since trade secret misappropriation law also comprises case law, attorneys will draw upon precedent cases via analogies and distinctions to strengthen their arguments that all necessary leaf issues of the trade secret law rules are fulfilled, or not, given the facts in the case, and that the desired outcome is warranted.

3.2 Cases and Factors

As a running example for this section, consider the DYNAMICS case[3]. As a brief thought experiment, the case's factors can be considered as a narrative when embedded into boilerplate text:

> **DYNAMICS (1980)**
> Plaintiff had developed product information and was marketing a product based on the information. Plaintiff's information was unique in that plaintiff was the only manufacturer making the product. Plaintiff took active measures to limit access to and distribution of its information. Plaintiff disclosed its information in a public forum.
>
> At some point, defendant obtained the product information. Defendant entered into a nondisclosure agreement with plaintiff. The nondisclosure agreement did not specify which information was to be treated as confidential.
>
> Eventually, defendant developed a competing product and commenced to sell it. Thereafter, plaintiff brought suit against defendant for trade secret misappropriation.

[3]Dynamics Research Corp. v. Analytic Sciences Corp., 9 Mass. App. 254, 400 N.E.2d 1274, 209 U.S.P.Q. 321 (1980).

This description communicates the case's tension. The plaintiff has problems arguing that he has made effective efforts to maintain the secrecy of the information since he has disclosed the information in a public forum. On the other hand, the information was unique and defendant signed a nondisclosure agreement, albeit a nonspecific one. Connecting these facts to the elements of the VJAP domain model, one can draw the following semantic connections:

- F15$_{\pi}$[unique-product] : The plaintiff's information being unique is relevant for *info-valuable*
- F6$_{\pi}$[security-measures] : The plaintiff having taken active measures to limit access to and distribution of its information is relevant for *maintain-secrecy*
- F4$_{\pi}$[agreed-not-to-disclose] : The defendant having entered into a nondisclosure agreement with plaintiff is relevant for *breach-of-confidentiality* and *maintain-secrecy*
- F5$_{\delta}$[agreement-not-specific] : The fact that the nondisclosure agreement was not specific is relevant for *breach-of-confidentiality*
- F27$_{\delta}$[disclosure-in-public-forum] : The plaintiff having disclosed its information in a public forum is relevant for *maintain-secrecy*

Figure 1: VJAP domain model of issues and factors

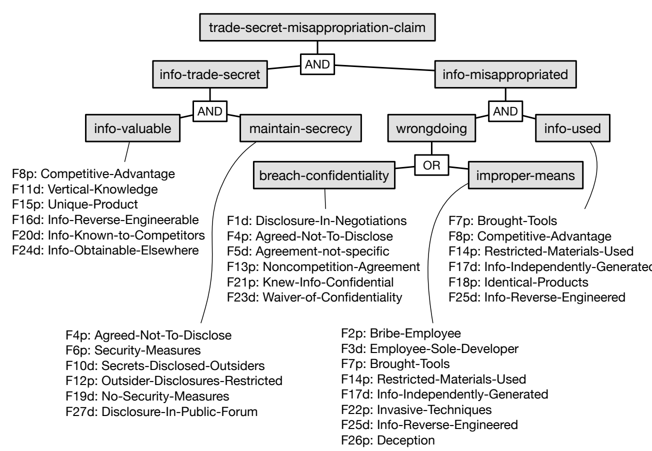

3.3 Leaf Issue Statuses and Presumptions

A formal representation of DYNAMICS is displayed in Fig. 2. One can recognize the plaintiff's dilemma in trying to establish *maintain-secrecy* and *breach-of-confidentiality* as both issues are *contested* by factors favoring the defendant. *info-valuable* , on the other hand, only has a plaintiff factor. No factors for *info-used* and *improper-means* are present in the case. In the VJAP model, if a case contains no factors for a leaf issue, we assume it can be considered *not raised* and one of two alternatives applies. For *info-valuable* , *maintain-secrecy* and *info-used* the issue is presumed to have been conceded by the defendant and won by the plaintiff. This is because the cases were annotated from written decision texts that may omit facts undisputed by the parties. In DYNAMICS, *info-used* is presumed conceded and won by the plaintiff.

In the alternative, *breach-of-confidentiality* or *improper-means* will not be considered conceded by the defendant, but as simply not raised. This is functionally equivalent to being won by the defendant since the plaintiff has the burden of proof. However, if both *breach-of-confidentiality* and *improper-means* are not raised in a case, they likely must have been undisputed and the defendant's counsel strategized to challenge the suit on the ground that the information is not a trade secret. In such a case, *wrongdoing* is considered conceded and won by the plaintiff.

Figure 2: DYNAMICS in the VJAP domain model

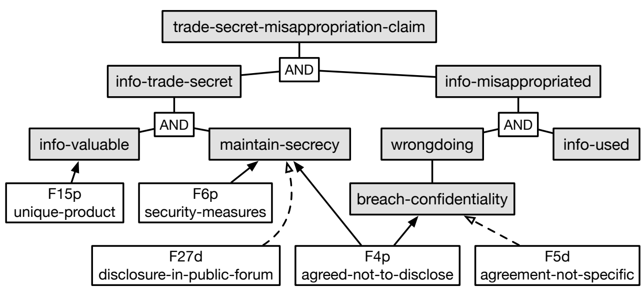

3.4 Relationship to IBP

VJAP's domain model is a variation of the domain model developed by Brüninghaus & Ashley for the issue-based prediction system (IBP) [12], with two modifications. First, in IBP, $\mathrm{F11}_{\delta}$ (a situation where the plaintiff's information was about customers and suppliers) was not associated with any issue. VJAP associates it with *info-valuable* because customer and supplier information arguably is not valuable information in the sense of the Restatement because it is deemed to be available from public sources. Second, in IBP, improper means alone established a misappropriation, even without the information being used. In VJAP, the Restatement rule is interpreted as the information needing to have been used for both improper means and breach of confidentiality

VJAP is also different from IBP in that it embeds the Restatement model into an argument generation system capable of value-based legal reasoning. Moreover, in IBP, so-called 'knockout factors' are used by the algorithm to discard cases that are counterexamples to the hypothesis that, when deciding a contested leaf issue, a set of plaintiff factors trumps a set of defendant factors. These factors are deemed to be so strongly in favor of the defendant that they prevent cases from always being considered useful counterexamples regarding any specific issue with which they are not associated. VJAP does not use such manually chosen knockout factors and, as will be explained, compensates through tradeoff argument schemes.

3.5 Trade Secret Domain Values

VJAP models four values underlying the trade secret domain. The most complex value is the *protection of the plaintiff's property interest.* Through innovative research, companies generate information which has intrinsic value because it can be monetized by means of developing a product based from it. It thus confers an advantage over competitors who do not have this information. The plaintiff also has an interest in the *protection of confidentiality* of certain transactions, focusing not the value of the information but the safety in commercial transactions provided by the plaintiff's ability to rely on confidential information to not be made public or used wrongfully. The general public has an interest in the law protecting *fair competition* practices by providing that wrongfully caused damage is subject to liability and, if needed, punitive damages. Finally, the general public has an interest in *public information being effectively usable* to enable innovation and healthy competition.

Effects of Factors on Values Similar to [6, 7] and follow up work, VJAP associates factors with values but goes beyond promotion and demotion. Every factor-value relationship is labeled as one of six possible effect types across three complementary pairs.

More vs. Less Legitimate Interests: Certain facts increase the legitimacy of a subject's claim that her interest is warranted to be protected. For example, $\mathrm{F15}_{\pi}$[unique-product] increases the legitimacy of the plaintiff's property interest because a unique product can lead to a monetarily quantifiable competitive advantage. On the other hand, $\mathrm{F20}_{\delta}$[info-known-to-competitors] decreases the legitimacy of the plaintiff's property interest because product development information known to competitors, if at all, grants a much weaker competitive advantage and it may be coherent with the law to not interfere by awarding a misappropriation claim.

*Waivers vs. Protections of Interests:*Certain actions amount to waiving protectable interests. For example, in a case where $\mathrm{F1}_{\delta}$[disclosure-in-negotiations] applies, the plaintiff has arguably waived his property interest in the product information by disclosing it in negotiations. Other actions qualify as active protections of ones interest, possibly leading to greater legal protection. For example, in a case where $\mathrm{F13}_{\pi}$[noncompetition-agreement] applies, the plaintiff has protected his confidentiality interest by entering into a noncompetition agreement with the defendant.

Interference vs. Non-interference of Interests: Interests can be interfered with, or not. If the defendant obtained the product information through deception ($\mathrm{F26}_{\pi}$[deception]), then the general public's interest in adherence to fair competitive practices has been violated. If the defendant discovered the information through reverse engineering ($\mathrm{F25}_{\delta}$[info-reverse-engineered]), then this is evidence that the plaintiff's property interest has not been violated.

4 EFFECT TRADEOFFS

4.1 Global Effect Tradeoff

Deciding a tradeoff in a case can be done without recourse to any legal rules. We define the **global effect weight** in favor of a party α in a case c as the sum over all favorable effects E^+ on applicable values V incurred by the factors F_c in the case.

$$ew(\alpha, c) = \sum_{v_j \in V} \sum_{e_k \in E} \sum_{f_l \in F_c} \theta(v_j, e_k, f_l)$$

The global effect tradeoff in the case is then a comparison of the global effect weight for each plaintiff and defendant. Fig. 3 illustrates the global tradeoff between factor effects in DYNAMICS. The left side of the diagram shows plaintiff factors whose effects on the values in the center are displayed through labeled edges. Every edge corresponds to a quantitative parameter $\theta(v_j, e_k, f_l)$ and the total weight favoring the plaintiff is the sum of those parameters. The right side of the diagram shows defendant factors and their effects on the values. The global tradeoff can be thought of as the balancing act between the left and right side of the diagram.

Figure 3: Diagram of global tradeoff in DYNAMICS.

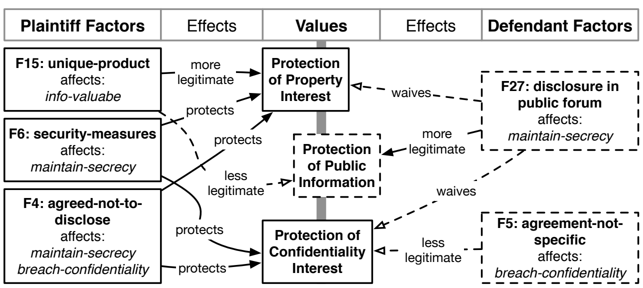

4.2 Scoping Effect Tradeoffs

VJAP improves upon the global tradeoff baseline by implementing scoped value judgments. One way to accomplish this is by referring to *local* effect tradeoffs inside an issue and then drawing upon precedent cases for support. The second approach is to engage in value effect balancing *inter-issue*, where a party's strong position on one issue makes up for its weakness in another. Again, such an argument can be supported by citing a precedent. The opponent can then attack the analogy, thereby undermining the value argument's persuasiveness, and the other party in turn can defend its analogy.

For tradeoff reasoning about a specific issue, the global tradeoff needs to be dissected into sub-tradeoffs (analogues to scoped value judgments), involving effect weight sums of smaller scope. We can quantify the **local effect weight** of a party α's argument in a given case c on a given issue i by restricting the effects to those from factors connected to the issue $F_c^i \subseteq F_c$ as follows.

$$ew(\alpha, c, i) = \sum_{v_j \in V} \sum_{e_k \in E} \sum_{f_l \in F_c^i} \theta(v_j, e_k, f_l)$$

From this we can compute the **local tradeoff confidence**:

$$\phi_{lto}(\alpha, c, i) = \frac{ew(\alpha, c, i)}{ew(\alpha, c, i) + ew(\overline{\alpha}, c, i)}$$

If the plaintiff has a local tradeoff with confidence x, then the defendant has an inverted tradeoff based on the same factors with confidence $1-x$. Fig. 4 shows the local tradeoff for *maintain-secrecy* between pro-plaintiff and pro-defendant effects. It is a subset of the global tradeoff in Fig. 3. It shows that plaintiff can argue that the effects of $F6_\pi$ and $F4_\pi$ outweigh those of $F27_\delta$.

Figure 4: Local tradeoff for *maintain-secrecy* in DYNAMICS

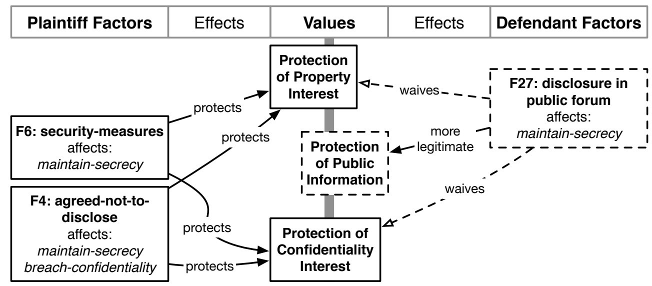

Consequently, **inter-issue tradeoff confidence** is derived from effect weights across a weak issue i_w and a strong issue i_s.

$$\phi_{iito}(\alpha, c, i_s, i_w) = \frac{ew(\alpha, c, i_s)}{ew(\alpha, c, i_s) + ew(\overline{\alpha}, c, i_w)}$$

Inter-issue tradeoffs also can be inverted for the opposing party, unless a side argues that its strong position on issue one outweighs its weak position on an issue where it has no factors at all.

Fig. 5 shows the inter-issue tradeoff for *breach-of-confidentiality* between pro-plaintiff and pro-defendant effects. It is also a partition (i.e. a subset) of the global tradeoff in Fig. 3. Plaintiff can compensate an arguably weak position ($F4_\pi$ vs. $F5_\delta$) by arguing that his strength on *info-valuable* (unchallenged $F15_\pi$) affords him a lower burden of persuasion on *breach-of-confidentiality*.

Figure 5: Inter-issue tradeoff for *breach-of-confidentiality* in DYNAMICS.

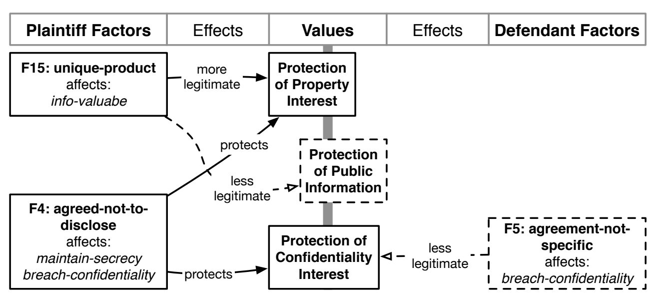

4.3 Confidence Propagation

VJAP generates an argument graph structure for every case it predicts using an exhaustive search through a backward chaining of argument schemes. Fig. 6 shows a pattern schema for the argument graphs that VJAP generates. It goes from arguments in the domain model at the top to deep arguments about leaf issues, tradeoffs, precedents, and analogy/distinction arguments between precedent and the case at bar. Once all graphs have been generated, VJAP takes a set of value effect weight parameters θ (compare definitions 16 and 17) and propagates a confidence measure across the graph towards the *trade-secret-misappropriation-claim* root statement, whose confidence is used in predicting the case outcome.

Confidence in a statement is aggregated in two ways, depending on the statement type. **Proportional confidence** accounts for all arguments about a statement. It is the sum of all pro-argument confidences over the sum of all arguments (pro and con). **Propmax confidence** is the confidence of the strongest pro-argument over the sum of the confidences of the strongest pro- and con-argument.

Figure 6: VJAP case argument graph schema

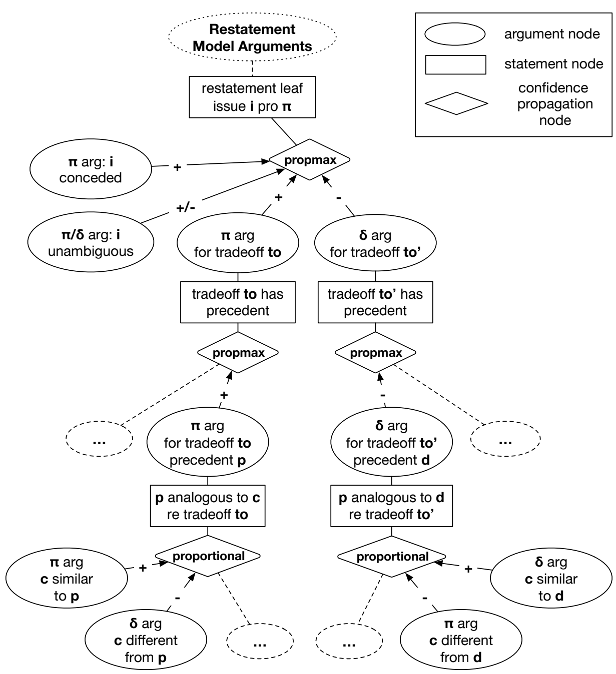

4.4 Argument Schemes

The VJAP model uses a set of schemes for basic arguments pro and con issues from sub-issues, conceded issues and undisputed issues, schemes about local and inter-issue tradeoffs, as well as corresponding precedents. Each scheme consists of a conclusion, a set of premises, a specialized function automatically computing the confidence in the conclusion from the confidence of the premises, and a verbalization function. The scheme set for local tradeoffs is given below. The set for inter-issue tradeoffs is analogous.

Argument for Issue from Local Tradeoff. A side argues that a given contested issue should be decided in its favor because a balancing of affected values related to the issue has to be decided in its favor.

Argument for Local Tradeoff from Precedent. A side argues for applying a local tradeoff to the case at bar and deciding a given issue in its favor because of a precedent case with the same tradeoff decided in its favor. Retrieved precedents may have different factors but are similar on the deeper level of their shared value effects. For example, in the case at bar defendant may have bribed one of plaintiff's employee's ($F2_\pi$[bribe-employee]) whereas in the precedent defendant used invasive techniques ($F22_\pi$[invasive-techniques]) to obtain the product information. Both cases share the deeper similarity that the defendant has violated the principle of fair competition. This can be thought of as pushing CATO's substitution argument move using higher-level concepts in the factor hierarchy [3] into the space of values and effects.

Argument for Local Tradeoff Precedent Analogy. A side argues that a case and a precedent which share the same local tradeoff of interest are sufficiently analogous that the precedent lends persuasive force to the application of the tradeoff in the case at bar.

Argument for Local Tradeoff Precedent Analogy from Supporting Surplus Factor in Current Case. In analogizing to a precedent, a party may point out a favorable factor that is present in the current case and not in the precedent but still part of the local tradeoff argued about in the analogy. The result is an *a fortiori* argument.

Argument against Local Tradeoff Precedent Analogy from Surplus Factor in Current Case. Analogies to precedents may be challenged by pointing out a factor in the current case disfavoring the side arguing for the analogy that is not part of the precedent, but should still be part of the tradeoff, thereby weakening the analogy.

Argument against Local Tradeoff Precedent Analogy from Surplus Factor in Precedent. Similarly to the previous scheme, the analogy can be challenged based on a surplus factor in the precedent that favors the challenging party and should be part of the tradeoff, thereby weakening the analogy.

4.5 VJAP Example Textual Argument

Below is an example inter-issue tradeoff argument generated by VJAP for the plaintiff on *maintain-secrecy* in ILG-INDUSTRIES:

> The plaintiff has taken efforts to maintain the secrecy of the information because disclosures to outsiders were subject to confidentiality restrictions. In fact, plaintiff's product information is so valuable that plaintiff must enjoy a lower standard to prove that the plaintiff has taken efforts to maintain the secrecy of the information because deciding otherwise would be inconsistent with the purposes underlying trade secret law.
>
> Specifically, regarding the value of the information, the product's uniqueness amounts to no scenario where usability of public information would be important and such a legitimate property interest regarding the value of the information that the confidentiality of outside disclosures alone must qualify as the maintenance of secrecy by the plaintiff despite the fact that plaintiff disclosed its product information to outsiders.
>
> A similar inter-issue tradeoff was made in ALLEN, which was decided for plaintiff. There, regarding the value of the information, the competitive advantage amounted to such a legitimate property interest and the product's uniqueness amounted to no scenario where usability of public information would be important and such a legitimate property interest that the maintenance of secrecy by the plaintiff was deemed sufficiently established despite the lack of strong evidence for plaintiff and fact that plaintiff had not adopted any security measures.
>
> Defendant may argue that ILG-INDUSTRIES is distinguishable from ALLEN because there access to the product information had saved time or expense, which is not the case in ILG-INDUSTRIES. Arguably, this made plaintiff stronger in the value tradeoff regarding the value of the information in ALLEN since plaintiff's property interest was more legitimate because of the competitive advantage.
>
> Defendant may argue that ILG-INDUSTRIES is distinguishable from ALLEN because there plaintiff had disclosed its product information to outsiders, which was not the case in ALLEN. Arguably, this makes plaintiff stronger in the value tradeoff regarding the value of the information in ILG-INDUSTRIES since plaintiff had waived his property interest and plaintiff had waived his confidentiality interest because of the disclosure to outsiders.

VJAP's texts contain redundant formulations and, when read over longer passages, can be recognized as having been generated by a computer program. However, one can imagine making the generation more natural by, for example, using varying formulations in a probabilistic grammar. Also, presumably, the arguments generated by VJAP do not necessarily correspond to the arguments made by the parties and the court in the actual cases that were the basis for the dataset, especially since VJAP does not use a semantic network representation like GREBE [11].

5 THE EXPERIMENT

5.1 Data

The original IBP case collection comprises 186 cases [5]. VJAP uses a subset comprising 121 (74 won by plaintiff, 47 by defendant) cases that have at least one factor for both plaintiff and defendant, thereby allowing balancing arguments. Seven cases turned out to have identical features (i.e. sets of factors). As very similar cases are not implausible in a legal setting, these seven cases were not removed from the data set. Also, a number of cases are sub-/supersets of other cases, leading to occasional *a fortiori* arguments.

5.2 The Experimental System

VJAP learns parameters from training cases and then predicts a set of test case using those parameters (see Fig. 7). For every case in the training set, the system generates an argument graph and then computes the confidence of the main issue statement regarding the plaintiff's claim by using the factor effect weight map parameters to calculate the confidence of the leaf-nodes of the graph. The argument schemes' and statements' confidence functions then propagate these values bottom up (or 'feed forward'). The case is predicted for the plaintiff if the confidence in the top level issue exceeds a threshold of .5, otherwise the defendant is predicted.

This process of top down graph construction, bottom up confidence propagation and winner prediction is done for every case in the training set and the overall prediction accuracy is calculated.[4] To learn optimal fact effect weight parameters, it happens in a loop during which the system iteratively searches for the optimal weight map using *simulated annealing* [22]. The training loop is run for a predefined number of iterations and the parameters (i.e. weights of factor effects on values) are adjusted at each iteration by replacing one random effect weight with a new random effect weight, thus generating a 'neighbor' in the space of possible weight maps. This new weight map's prediction accuracy is then computed through confidence propagation across all cases. If the neighboring weight map is better or equally good, it replaces the current weight map.[5]

At the end of the annealing process, the best found weight map is kept for testing. There, the test cases are predicted in the same construction-propagation-prediction fashion by using the trained effect weight parameters without any more optimization. The final accuracy is the number of correctly predicted test cases over the number of all test case predictions. Argument graphs for test cases can only be constructed using the training cases as precedents.

Figure 7: VJAP's system architecture

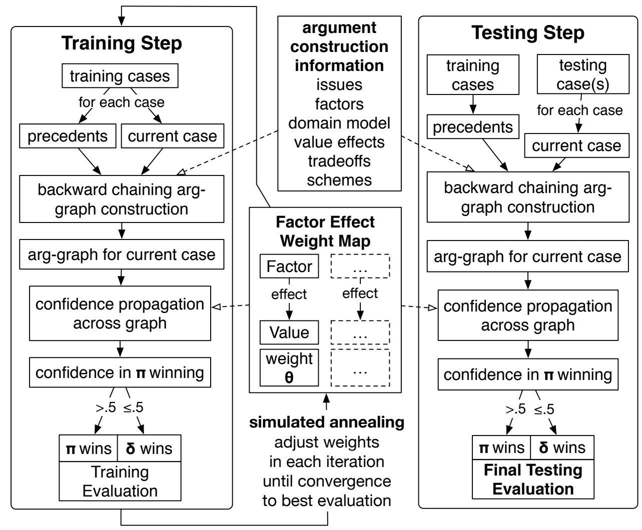

5.3 Experimental Conditions

Baselines. As baselines for evaluating the VJAP prediction performance, we use the following prediction system configurations.[6]
major-label: The simplest baseline always predicted the label of the majority of the cases, i.e. plaintiff wins the case.
Naive Bayes: A simple Naive Bayes classifier which performed second-best in an earlier evaluation of IBP [12].
Decision Tree: A C4.5 decision tree learner [25] also used in [12], which we chose over SVMs because of the small dataset.
IBP: A reimplementation of IBP's prediction algorithm given its assumptions regarding the domain model and factor associations.[7]
IBP-noEE: IBP without the functionality of explaining away counterexamples using knockout factors, and abstaining instead. This is to evaluate IBP's prediction algorithm without the benefit of knowing the impact of certain factors beforehand (see section 3.4).

VJAP Conditions. The following experimental configurations of the VJAP system were tested for their predictive performance:
Global Tradeoff Weight (GTW): Whichever side has the highest global effect weight wins, with plaintiff winning ties.
VJAP-full: Cases are predicted using the full argument graph including local and inter-issue tradeoff-based arguments.
VJAP-local: As VJAP-full, but restricted to local tradeoff schemes.
VJAP-no-precedent: As VJAP-full, but without precedent-based schemes, i.e. tradeoffs are resolved using only effect weights.
VJAP-timeline: As VJAP-full, but cases may only be argued with precedents that have been decided at least one year before the year

[4]Graphs are cached during experiment execution for better performance.
[5]The system may replace the current map with a worse one with a small probability computed (using a 'cooling schedule') from the system's 'temperature', which is a function of the remaining and total number of cycles in the annealing process. Such occasional 'bad moves' make the search less likely to get stuck in local optima in the space of possible weight map parameters.

[6]All standard machine learning models were trained using Weka 3.6.13 [21]
[7]It was reimplemented to the best of the author's ability from available documentation.

of the case at bar. This resembles real legal reasoning where cases form a system (or 'theory' [7, 23]) of the domain over time.

5.4 Experiment Configuration

Experimental conditions and baseline models were evaluated in a leave-one-out experiment and a 5-fold cross validation. The system first learns factor effect parameters using training cases and then predicts the test cases in single execution of the two steps shown in Fig. 7. The accuracy measure then reflects the system's performance on the test cases. In leave-one-out, the two steps are executed 121 times; each time another case becomes the 'single case test set' and the remaining 120 cases function as the training set. In 5-fold cross validation, the 121 cases were randomly assigned to 5 sets. In each fold, a different set of cases is used as the test set. For all VJAP and IBP models, the same 5 cross-folds were used.[8]

Unless otherwise specified, all VJAP experiments were run using simulated annealing with a linear cooling schedule [22] on 300 iterations, for the first 25% of which a full new random weight map was generated. For the remaining iterations the last best random weight map was modified by replacing a random weight with a new random value between 0.00 and 0.99 in 0.01 increments.

5.5 Results

Prediction Performance Evaluation. Table 1 presents the prediction performance results. All measures are accuracy, i.e. number of correct predictions over the total number of predictions made. [9]

VJAP vs. Machine Learning Baseline. Cross-validation results are close to the leave-one-out experiment for all conditions. VJAP-full and VJAP-timeline perform slightly worse, likely due to less training data in each cross fold than in leave-one-out. Naive bayes performs better in cross-validation (.851) than in leave-one-out (.843). C4.5 decision trees perform slightly worse on leave-one-out (.777) yet are second best (.835) in cross-validation. An explanation is that five folds in cross-validation reduce overfit and allow the model to perform better overall. Due to VJAP-full and VJAP-timeline doing slightly worse than in leave-one-out, naive bayes and decision tree baselines perform best during cross validation, albeit by only a small margin above VJAP-timeline.

VJAP vs. IBP. IBP performs at .81 in the leave-one-out.[10] In cross-validation, the model degrades to .725 because IBP's prediction algorithm uses the cases directly as points of reference without any separate trained parameters. VJAP works with its full set of parameters even if less training data is available.

VJAP-local is conceptually close to IBP because its local tradeoff arguments are functionally similar to IBP's issue analysis and theory testing. However, IBP usually performs better than VJAP-local, which may be due to IBP's limited ability to reason across issues

Table 1: Results table of predictive accuracy

model	LOO	train	5-fold	train	fold-SD
GTW	.802	.84	.805	.845	.059
VJAP-full	.793	.828	.779	.837	.046
VJAP-local	.694	.717	.691	.725	.095
VJAP-no-precedent	.711	.727	.715	.73	.097
VJAP-timeline	.843	.854	.821	.862	.079
major-label	.612	*n/a*	.612	*n/a*	*n/a*
naive-bayes	.843	*n/a*	.851	*n/a*	*n/a*
decision tree	.777	*n/a*	.835	*n/a*	*n/a*
IBP	.81	*n/a*	.725	*n/a*	.112
IBP-noEE	.587	*n/a*	.562	*n/a*	.14

using the hand-flagged knockout factors, which VJAP-local does not have. When IBP is deprived of this knockout factor functionality in the IBP-noEE condition, its performance drops below that of the major label baseline. The best performing experimental condition, VJAP-timeline, consistently outperforms IBP.

Comparison Among VJAP Models. In leave-one-out, GTW performs surprisingly well at .802, which is about the same as VJAP-full. Also, GTW's performance does not degrade in the cross-validation condition, suggesting that the argument-based VJAP models are more dependent on having access to a larger pool of training cases and citable precedents for argument construction. VJAP-full performs at .793 in leave-one-out and decreases slightly in cross validation. VJAP-local (.694) and VJAP-no-precedent (.711) perform significantly worse and do not degrade during cross-validation. VJAP-local does not have access to inter-issue tradeoffs and performs badly, which confirms that reasoning in the logically connected domain model benefits from recourse to some plus/minus tradeoff interactions across leaf issues.

This is strong evidence for VJAP's systemic assumption that requirements of legal rules are not semantically separated from each other. Rather, they are interdependent regarding the question of whether one specific requirements is deemed fulfilled given the specific facts of the case, the fulfillment of its 'semantically connected requirements', and the raised purposive concerns. Legal rules discretize and structure the legal decision process into smaller issues which, however, cannot be viewed in isolation.

VJAP-full performing better than VJAP-no-precedent suggests that computing confidence in tradeoff arguments from the effect weight parameters alone (i.e. without the influence of precedent arguments) fails to account for the normative force of the principle that cases, when litigated, are decided with recourse to prior case law. Hence, adherence to precedent forms a fifth quasi-value next to the four values VJAP models explicitly and, when taken into account through precedent argument schemes, contributes to the higher predictive power of VJAP-full and VJAP-timeline.

VJAP-timeline performs best of all VJAP models in the leave-one-out (.843, on par with naive bayes) and cross validation conditions (.821, slightly lower than the highest machine learning baseline). The plus in accuracy comes at the cost of a higher standard deviation across the folds (.079) when compared to VJAP-full (.046).

[8] Weka's cross validation function was used for the naive bayes and decision tree baselines, resulting in different random foldings for these two conditions.

[9] Average training accuracy and standard deviation across folds are not given for naive bayes and decision trees because the standalone version of Weka does not provide these quantities by default without a manual implementation. Training error for IBP models is not available because IBP does not have a training stage but rather uses its repository of cases directly in the prediction process.

[10] It should be noted that IBP's earlier experiments in [12] were done on the full 185 case dataset and reported a .91 accuracy in the leave-one-out condition. This experiment only uses a subset of 121 cases that contain factors for both sides.

This suggests that the restriction to temporally plausible precedent arguments makes the system more sensitive to folding because it reduces the scope of possible precedents for cases even further. VJAP-full's unrestricted ability use all available training cases as precedents seems to produce a more stable prediction performance.

Decision Threshold Variation. The .5 decision threshold was deliberately set assuming that simulated annealing is able to adapt the weights accordingly, and that the weight maps created across runs are comparable. Arguably, it should be another parameter learned from data. To assess the impact of assuming a static threshold, VJAP-timeline was run on the full dataset with thresholds from .05 to .95 in .05 increments. Fig. 8 shows how different thresholds affect prediction accuracy for cases won by plaintiff, cases won by defendant, and all cases. Performance peaks at around .45, confirming that choosing a midway threshold was reasonable.

Figure 8: prediction accuracies for different thresholds

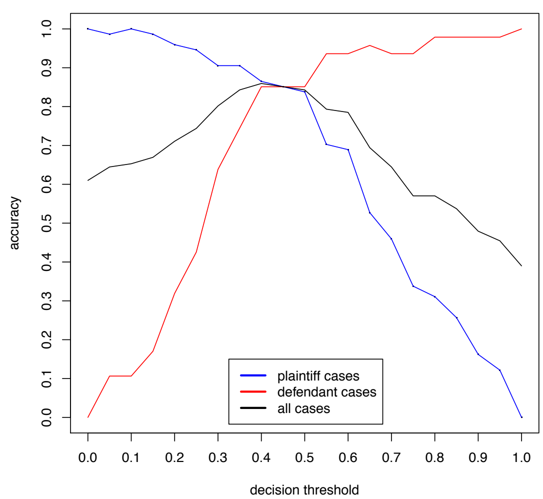

Citation Graph Analysis. VJAP-timeline predicts case outcomes using temporally plausible precedents. Using a trained weight map, the system generated the strongest plaintiff and defendant argument for all five leaf issues of the domain model. If the strongest argument for a given side is a tradeoff-based argument that involved a precedent analogy, then that precedent is considered 'cited'. These citation links exhibit interesting patterns, but do not necessarily reflect the real citations in the opinions of these cases.

Influential cases tend to have many factors, contain many possible tradeoffs and, accordingly, can be cited by subsequent cases with fewer factors. These cases, in turn, tend to be less cited because they are less complex. Fig. 9 shows an excerpt of the citation graph where two complex cases (STRUCTURAL-DYNAMICS and MOTOROLA) are cited by four simpler cases which in turn are not cited by subsequent cases (they do not have incoming arrows). Specialized 'lines of cases' are also formed. Fig. 10 shows a line of cases for the defendant, which traverses CORROON, FREEDONIA and OPTIC-GRAPHICS with the latter three having identical factors but having been decided at different points in time. FREEDONIA

Figure 9: Citation graph excerpt for complex/simple cases; Blue/red arrows are plaintiff/defendant citation links

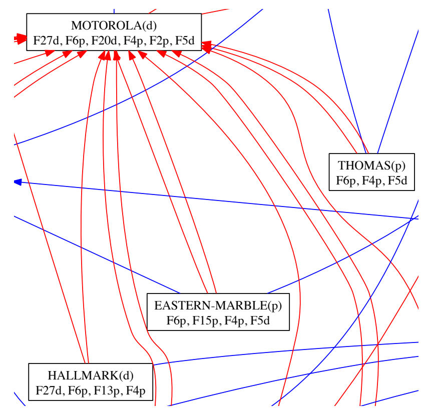

Figure 10: Citation graph excerpt showing line of cases

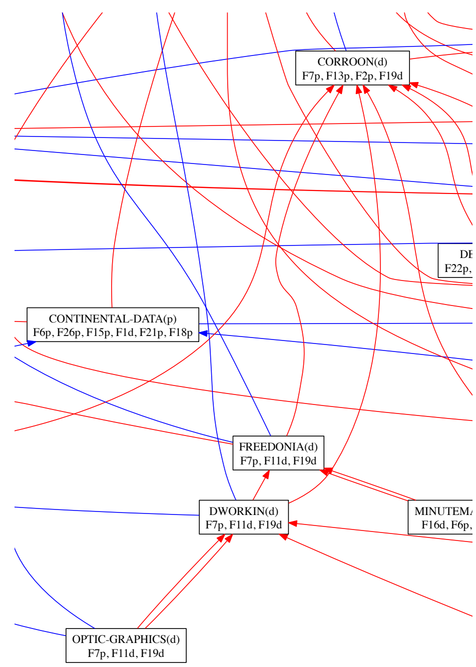

cites CORROON for one argument. DWORKIN is identical to FREEDONIA, but relies on both FREEDONIA and the more influential CORROON. OPTIC-GRAPHICS then relies entirely on DWORKIN.

Incorrectly Predicted Cases. Some of VJAP's predictions of the winning party are incorrect. For example, early cases cannot cite any precedents for supporting tradeoff arguments which then miss out on the added confidence boost of a precedent backing. The same is true of cases with very few factors that are so specific that VJAP cannot find any precedent. Also, cases won by plaintiff with issues that exclusively have defendant-factors can only be predicted correctly by means of inter-issue tradeoffs and strong precedent to acquire enough confidence for the plaintiff's argument.

6 RELATED WORK

HYPO's dimensions [4] correlate to 'degrees of support' for a party's argument but without an explicit representation of values. CATO [3] simplified dimensions into binary factors and interrelated them in a hierarchy, allowing for more sophisticated arguments, although again without making values explicit.

Branting's GREBE [11] was a sophisticated hybrid rule- and case-based reasoning system which relied on manually created common sense knowledge and uses structure mapping in a semantic network representation to retrieve and compare cases. While GREBE's goes a long way in representing the intricacies of its domain in a deep model, it did not, to the author's best knowledge, invoke purposes or values as justifications in the arguments it produced.

Bench-Capon & Sartor have cast case-based reasoning as a form of *theory construction* [7]. They decide cases using rules, conflicts among which are resolved using a hierarchy of values they promote or demote. Preferences among values are themselves propositions that may be argued about in, for example, hierarchical argumentation frameworks [24]. Chorley & Bench-Capon replaced the qualitative ordering of values with a quantitative weight for values and factors, and reported results on a publicly available subset of IBP cases [15]. They further structure values through degrees of promotion and demotion [14]. The AGATHA system [13] operationalized argument-like theory constructors in [7] in an adversarial-like argument game. Greenwood et al. [20] models CATO-like case-based legal reasoning as an instance of argumentative persuasion focusing on the effects of an action on values in a situation towards a certain goal. Later work [9] uses value-based, contextual preferences between actions for practical reasoning, but does not apply it to case-based reasoning or outcome prediction. VJAP does not use abstract, qualitative preference relations between values as in [7, 8, 20], or derivatives thereof. Instead, it balances effects of factors on values in a given case as well as analogizes and distinguishes precedents containing similar tradeoffs. It further captures temporal dynamics in its model of precedent-based reasoning.

7 LIMITATIONS & FUTURE WORK

Once more annotated cases are available, VJAP's prediction performance has to be subjected to a rigorous empirical evaluation on a larger dataset satisfying statistical requirements. Also, the VJAP repository of analogizing and distinguishing moves is still limited. The prior explorations of emphasizing and downplaying of similarities and differences, as well as hypothetical reasoning, in the VJF [18, 19] provide a starting point for additional argument schemes. Future work can further explore more fine-grained knowledge representations of case facts, similar to that of GREBE [11]. Like GREBE, VJAP could eventually be evaluated by a human assessment of its generated arguments.

8 CONCLUSIONS

VJAP goes beyond the scope of prior systems in AI&Law because of the way it accounts for values, its ability to learn quantitative effect weights from prior cases, its use of argument schemes to generate case-based legal arguments arguments, as well as its assessment and use of these arguments to predict case outcomes. Results show that its prediction performance is comparable with that of common machine learning methods when taking into account the chronology of its case-base, but it also generates realistic legal arguments. A comparison of different experimental configurations reveals empirical evidence that implementing tradeoffs and augmenting them with precedent-based argumentation produces better prediction results, and that parallel antecedents of legal rules interact by means of balancing in plus-minus configurations.

REFERENCES

[1] Alexy, R., On Balancing and Subsumption: A Structural Comparison, Ratio Juris, 2003, 16, 433-49

[2] Ashley, K. D., Evaluating the uses of Values in A model of Legal Reasoning with Cases Incorporating Theories and Values, From Knowledge Representation to Argumentation in AI, Law and Policy Making, Festschrift in Honour of Trevor Bench-Capon, 2013

[3] Aleven, V., Using background knowledge in case-based legal reasoning: a computational model and an intelligent learning environment, Artificial Intelligence, Elsevier, 2003, 150, 183-237

[4] Ashley, K. D., Modeling legal arguments: Reasoning with cases and hypotheticals, MIT press, 1991

[5] Ashley, K. D. & Brüninghaus, S., Computer models for legal prediction, Jurimetrics, 2006, 309-352

[6] Atkinson, K. & Bench-Capon, T., Legal case-based reasoning as practical reasoning, Artificial Intelligence and Law, Springer, 2005, 13, 93-131

[7] Bench-Capon, T. & Sartor, G. A model of legal reasoning with cases incorporating theories and values, Artificial Intelligence, 2003, 150, 97-143

[8] Bench-Capon, T., Persuasion in practical argument using value-based argumentation frameworks, Journal of Logic and Computation, 2003, 13

[9] Bench-Capon, T. & Atkinson, K., Action-state semantics for practical reasoning, AAAI Fall Symposium, Technical Report FS-09-06, 2009, 8-13

[10] Berman, D. H. & Hafner, C. D., Representing teleological structure in case-based legal reasoning: the missing link, Proceedings ICAIL 1993, ACM, 50-59

[11] Branting, L. K., Building explanations from rules and structured cases, International journal of man-machine studies, Elsevier, 1991, 34, 797-837

[12] Brüninghaus, S. & Ashley, K. D., Predicting outcomes of case based legal arguments, Proceedings ICAIL 2003, ACM, 233-242

[13] Chorley, A. & Bench-Capon, T. AGATHA: Using heuristic search to automate the construction of case law theories, Art. Int. and Law, 2005, 13, 9-51

[14] Chorley, A. & Bench-Capon, T., An empirical investigation of reasoning with legal cases through theory construction and application, Artificial Intelligence and Law, 2005, 13, 323-371

[15] Chorley, A. & Bench-Capon, T., Reasoning with legal cases as theory construction: Some experimental results, Proceedings JURIX 2003, IOS Press, 2003

[16] Feteris, E. T., The role of pragmatic argumentation referring to consequences, goals, and values in the justification of judicial decisions, Proc. ISSA 2014, 414-425

[17] Grabmair, M., Modeling Purposive Legal Argumentation and Case Outcome Prediction using Argument Schemes in the Value Judgment Formalism, Doctoral Dissertation, University of Pittsburgh, 2016

[18] Grabmair, M. & Ashley, K. D., Argumentation with Value Judgments - An Example of Hypothetical Reasoning, Proceedings JURIX 2010, 67-76

[19] Grabmair, M. & Ashley, K. D., Facilitating case comparison using value judgments and intermediate legal concepts, Proceedings ICAIL 2011, ACM, 161-170

[20] Greenwood, K.; Bench-Capon, T. & McBurney, P., Towards a computational account of persuasion in law, Proceedings ICAIL 2003, 22-31

[21] Frank, E., Hall, M., Witten, I.H., Weka-a machine learning workbench for data mining, Online Appendix for "Data Mining: Practical Machine Learning Tools and Techniques", Morgan Kaufmann, 2016

[22] Kirkpatrick, S., Gelatt, C. D., & Vecchi, M. P., Optimization by simulated annealing, Science, 1983, 220(4598), 671-680

[23] McCarty, L. T., An implementation of Eisner v. Macomber, Proceedings ICAIL 1995, ACM, 276-286

[24] Modgil, S., Value Based Argumentation in Hierarchical Argumentation Frameworks, Proceeding COMMA 2006, IOS Press, 297-308

[25] Quinlan, J. R., C4. 5: programs for machine learning, Elsevier, 2014