

An Abductive Theory of Legal Issues

Thomas F. Gordon*

German National Research Center for Computer Science (GMD)
Institute for Applied Information Technology
Expert Systems Research Group
Sankt Augustin, Federal Republic of Germany
thomas@gmdzi.gmd.de

December 12, 1990

Abstract

A normative theory of legal issues, argument moves and clear cases is presented in which abduction, rather than deduction, is of central importance. The theory is a refinement of Fiedler's constructive view of legal reasoning. Like legal positivism, the theory elaborates the concept of a clear case, but here clearness is defined relative to a set of competing interpretations of the law, rather than a single consistent set of "valid" rules. A computational model for the theory is described, which uses an ATMS reason maintenance system to implement abduction. Finally, the theory is compared with Anne Gardner's program for spotting issues in offer and acceptance law school examination questions.

1 Introduction

This paper describes a normative theory and computational model of certain aspects of legal argumentation. That is, it is not a descriptive theory of actual legal reasoning, but proposes standards against which actual reasoning can be evaluated. Although the computational model may have relevance for legal expert systems, its main purpose is to allow the theory to be tested by simulating legal arguments.

Aspects of legal reasoning addressed in this theory are issue spotting, certain kinds of argumentation moves, and the notion of clear cases. The approach to issue spotting is compared with Gardner (1987). The theory was first published in Gordon (1989), but is presented here in a new way, building upon an abstract theory of abduction, instead of directly upon the implementation using de Kleer's ATMS reason maintenance system.

The paper is organized around five main sections, excluding this introduction and the conclusion: Section 2 provides the necessary jurisprudential orientation, comparing legal positivism with Fiedler's constructive theory of legal reasoning. Section 3 discusses abduction, which plays a central role in the theory of issues, argumentation and clear cases which follows, in section 4. Section 5 describes briefly the the computational model implementing the theory. Finally, section 6 is a comparison of my theory with Gardner's program for spotting issues in offer and acceptance examination questions.

*The research reported here was conducted as part of the Assisting Computer project (AC) of GMD's Institute for Applied Information Technology. I would like to thank Gerhard Brewka, Joachim Hertzberg and Ulrich Junker for their collegiality and L. Thorne McCarty and Wolfgang Bibel for their careful editing.

2 Jurisprudential Background

Fiedler (1985) presented a vision of a kind of Computer-Aided Design (CAD) system for “drafting” legal arguments. In that paper, he sketched out a theory of jurisprudence which I will refer to as a *constructive* or *abductive* theory of legal reasoning. My own abductive theory of legal argumentation, presented here, is a step towards a precise mathematical and computational refinement of Fiedler’s rather abstract presentation.

Before describing Fiedler’s theory in more detail, let us briefly discuss *legal positivism*, the theory of legal reasoning most widely accepted today, at least by practicing lawyers in common law jurisdictions, Hart (1961). Summarizing Dworkin (1977a), central properties of Hart’s particular version of legal positivism include:

1. The law consists of a set of valid rules, which can be identified by applying a fundamental secondary *rule of recognition*;
2. The valid rules are incomplete. Some cases, the *clear cases*, are decidable by applying the rules; the others require the exercise of *judicial discretion*.
3. Legal obligations arise only out of valid legal rules. In a hard case, a party may be held liable for an obligation which did not exist at the time of the events of the case.

Legal positivism’s recognition that the rules have gaps distinguishes it from mechanical jurisprudence. However, there is a trace of mechanical jurisprudence left in legal positivism’s notion of clear cases: once the relevant valid rules have been identified and the case has been discovered to be clear, its decision follows deductively by a mechanical application of the rules to the facts.

This brief account of positivism obviously cannot do justice to all its versions by various authors. Also, Hart himself acknowledged that there are limits to his rule of recognition, by admitting, for example, that there is no “. . . authoritative or uniquely correct formulation of any rule to be extracted from cases.” Hart (1961), pg. 131. But he goes on to claim that this is not usually a problem: “. . . there is often very general agreement . . . that a given formulation is adequate”.

When speaking of positivism here, then, I mean a reconstruction of positivism in its simplest, pure form. It is supposed that the rule of recognition can be applied to the primary sources, such as statutes and cases, to construct a single, consistent theory of the valid rules of law, represented as a set of propositions in some first-order language. For each open-textured predicate, such as “vehicle”, it is supposed that there is a rule in the theory expressing its “core of certainty”, such as $\text{car} \rightarrow \text{vehicle}$. Let Γ be such a theory of the law. Then, given a set of propositions Θ , describing the facts of the case, and some goal ϕ , the case is *clear* just when $\Gamma \cup \Theta \models \phi$ or $\Gamma \cup \Theta \models \neg\phi$. Otherwise, the case raises hard questions requiring the exercise of judicial discretion.

Legal positivism was explicitly adopted as the philosophical basis for at least two AI and Law projects. In Susskind (1987), legal positivism is used to defend purely deductive legal expert systems. His main point is that, despite the various sources of uncertainty in legal reasoning, because of the “core of certainty” of legal rules postulated by legal positivism in all its versions, deductive expert systems can be useful tools for *lawyers*, who can be trusted to respect the limits of the system and know when other methods are required. On the other hand, Gardner (1987) directly addressed the problem of identifying the *hard questions* raised in offer and acceptance law school examination questions. Unlike Susskind, who was principally interested in tools for lawyers, Gardner’s interest was developing a computational model of legal reasoning capable of autonomously spotting hard legal issues. The easy questions were still resolved deductively.¹

One of legal positivism’s most influential critics, Ronald Dworkin, argues that *principles* as well as rules constrain judicial discretion, Dworkin (1977b). Dworkin might find some cases to be

¹Gardner did not restrict her view of positivism to the pure form adopted here. Her program reasoned with multiple interpretations of cases, and considered a hard question to have been raised if a predicate defined by competing rules was relevant.

clear, because principles converge to a particular decision, where positivism would consider the law to contain a gap requiring the exercise of discretion.

Within the field law and AI, there had been a debate about the utility of rule-based legal expert systems, which was considered to hinge on whether or not there are clear cases, as supposed by legal positivism. Leith took the position that there are no clear cases, giving several examples of apparently clear cases being decided in unexpected ways, Leith (1985). In Susskind (1987), pg. 238, Richard Susskind argued that Leith misunderstands legal positivism, because he speaks of “clear rules” rather than clear cases.²

Leith’s criticism of positivism, at least in the pure form considered here, is well-founded. He correctly points out that positivism does not adequately account for the process of *destructively* modifying rules during the decision of cases. Rather, positivism supposes that the set of rules is merely *extended* with additional rules for the decisions of what had been hard cases, thus reducing the *penumbra of doubt* in future cases.

In section 4.3, a new theory of clear cases is proposed which avoids these problems by defining clearness relative to a fixed set of competing interpretations of the law. Briefly, a case is clear just if all known interpretations of the law lead to the same result. Future decisions, extending the set of interpretations, can retroactively cause a case to be clear which had been hard, or vice versa.

Now, let us return to Fiedler’s constructive theory of legal reasoning, to see how it differs from legal positivism. The theory starts by recognizing that legal reasoning, in its full breadth and richness, consists of a variety of reasoning tasks, such as:

- Identifying and interpreting relevant legal texts, such as cases and statutes;
- Constructing alternative views of the facts;
- Designing legal arguments.
- Managing dependencies between versions of facts, alternative interpretations of legal texts, arguments and inferences;
- Selecting “best” arguments; and
- Documenting chosen arguments persuasively.

This list is of course not exhaustive. Also, the ordering of the tasks is not meant to suggest that the tasks are carried out in the order given. On the contrary, it is supposed these tasks are performed *opportunistically*, Hayes-Roth and Hayes-Roth (1979). That is, a lawyer will jump from one task to another to take advantage of opportunities as they present themselves.

Clearly, when legal reasoning is characterized in this broad, diverse way, it is difficult to see it as being primarily a deductive task, although deduction will play an important, if subordinate role, in performing some of these tasks. Rather, according to Fiedler (1985):

... the task of the judge essentially includes the choice, shaping and logical construction of the appropriate legal rules as well as the pertinent statements of facts in mutual interdependence. It is true that the resulting fabric of the judgment and its reasons has the function of deductively connecting the decision to the rules of law and the facts of the case. Nevertheless the process of decision-making is not reduced to the application of deductive logic to given premises, but essentially consists in constructing a logical fabric, which at the same time has the qualities of an adequate model and a stringent deduction. In terms of modern methodology, judicial decision-making will have to be qualified as a process of model-construction or “modeling”.

²However, at least one reknowned legal positivist has used the term “clear rule”, perhaps as a shorthand way of suggesting that the ratio of “core of certainty” to “penumbra of doubt” is quite large, MacCormick (1978), pg. 198.

Thus, the role of deduction in Fiedler’s conception of legal reasoning, which coincides with the view in Alchourron and Bulygin (1989), is principally one of *justifying* decisions, after they have been made, at least tentatively, rather than providing a method for reaching or discovering decisions. That is, the decision must follow logically from the rules of law and the facts of the case, but the rules used are not present in the argument from the beginning; they are selected from the body of precedents and “shaped”, if necessary, by the lawyer during his construction of an argument justifying the decision. The rules are shaped by interpreting, and reinterpreting, the various sources of law, such as reported cases, and the facts arise out of the interpretation and discovery of evidence. The choice of terminology for stating the facts and the rules are mutually dependent on one another. The law cannot be formulated without considering the facts, and the facts cannot be stated without considering the applicable law. Logical deduction is principally a *constraint* on the structure of the resulting argument. In AI terms, legal reasoning according to Fiedler’s constructive theory can be viewed as search through the space of interpretations of legal texts and facts of the case.³

Fiedler’s theory does not elaborate how legal texts can or should be interpreted. It only asserts that arguments are constructed through the “choice, shaping and logical construction of the appropriate legal rules ...”. The theory supposes that a set of established, competing interpretations exists before one begins to construct arguments for some new case. This is the case, for example, when courts in different jurisdictions have decided some question of law in contradictory ways. Which interpretations of primary legal sources are of interest? Although every interpretation, from whatever source, may be of interest when trying to make a creative, new argument, it is the established precedents contained in the published decisions of the courts which are of primary importance, as it is surely a reasonable strategy to begin reasoning with precedents before attempting to break new legal ground. Thus, the interpretations of law of interest are of two basic kinds: published precedent and new creative interpretations.

This constructive view of legal reasoning differs from legal positivism primarily in that it does not suppose there to be a “rule of recognition” for identifying a consistent, unique set of legal rules. Although positivism accepts that the set of “valid” rules will be incomplete, the constructive view goes further; it acknowledges that the interpretations to be considered will often be inconsistent. Positivism’s “rule of recognition” is replaced by a weaker test, one that permits alternative reasonable interpretations of the primary sources.

3 Abduction

In Gordon (1989), my theory of issue spotting was described directly in terms of the implementation, which uses the ATMS presented in de Kleer (1986). Here, I recast the theory more abstractly in terms of abduction and then show, in the section on the computational model for this theory, how it may be implemented using the ATMS.

But just what is abduction? Until recently, the term does not seem to have attracted very much attention in the fields of Artificial Intelligence and Philosophical Logic. It appears neither in the index to the four volume Handbook of Philosophical Logic, Gabbay and Guenther (1983), nor in the index of the Kneales’ 800 page history of logic, Kneale and Kneale (1962). In the AI literature, there is a chapter on abduction in the Introduction to Artificial Intelligence, Charniak and McDermott (1985), but they do not tell us where the word was first used in its logical sense. Surprisingly, the Logical Foundations of Artificial Intelligence also does not include a discussion of the topic, Genesereth and Nilsson (1987).

One of the few definitions of abduction I have seen appears in the Encyclopedia of Philosophy, Edwards (1972), vol. 5, pg.57:

- (1) A syllogism whose major premise is known to be true but whose minor premise is merely probable.
- (2) C.S. Pierce’s name for the type of reasoning that yields from a

³Edwina Rissland described the task performed by Gardner’s program in the same terms, Rissland (1988).

given set of facts an explanatory hypothesis for them.

The example chosen by Charniak and McDermott has us conclude that Jack is drunk from seeing that he is not walking straight. Here, we “jump” to the conclusion by applying common sense knowledge about drunkenness. In propositional logic, we could express this reasoning as follows:

$$\frac{d(j) \rightarrow \neg ws(j), \quad \neg ws(j)}{d(j)}$$

Here, $d(j)$ and $ws(j)$ are intended to mean *Jack is drunk* and *Jack walks straight* respectively. As is readily apparent, abduction is seen here to be a kind of *plausible* inference method. It may be that Jack is not walking straight for some other reason. However, I find this particular formulation of abduction as a kind of inference rule to be somewhat confusing, as it fails to emphasize the difference between abduction and deduction. Rather, I prefer to view abduction as a form of theory construction. Given some general knowledge about the world, $\{d(j) \rightarrow \neg ws(j)\}$, and a set of hypotheses, $\{d(j), \dots\}$, try to find a subset of the hypotheses from which, together with the general knowledge, the observation of interest, $\neg ws(j)$ here, is a logical consequence in the standard sense:

$$\{d(j)\} \cup \{d(j) \rightarrow \neg ws(j)\} \models \neg ws(j)$$

As in the example, abduction is usually seen as a method for generating explanations for observations. However, as Charniak and McDermott point out, it is not completely satisfactory for this purpose, as material implication is being used to represent causality. This works fine in the example, but α is usually not a cause of β in sentences of the form $\alpha \rightarrow \beta$. To use another example from Charniak and McDermott, it may be that all patients in a cancer ward have cancer, but one would not want to accept their being in the ward as an explanation for their cancer.

Let us adopt a more abstract view of abduction which avoids these problems with causality and broadens abduction’s range of applicability. In deduction, we ask what are the logical consequences of some given set of sentences, or whether some particular sentence is a consequence of the set. In abduction, we are given a set of sentences, Γ , which is not expected to be consistent, and some sentence of interest, ϕ . The task is to find the smallest, consistent subsets of Γ that entail ϕ .⁴

Now, such a conception of abduction simply side-steps the problem of causal explanations of observations, as it is completely divorced from this particular application. If ϕ is some observation to be explained, and Γ includes the alternative explanatory hypotheses, it may still be that ϕ is entailed by certain subsets of Γ which we would be unwilling to accept as explanations. We can consider logical consequence here to be a *constraint* or necessary condition of an explanation, but it alone is not sufficient. Other criteria are required to filter out the nonexplanatory theories.

More importantly for our purposes here, this abstract interpretation of abduction is more suitable for use in a theory of legal reasoning. Legal decisions are not observed but made, and it is not our purpose to explain judgments but to justify them.

As in the case of explaining observed events, I suppose it will not be enough to merely show that the legal decision is entailed by some consistent theory. A complete theory of “technically correct legal arguments”, to use Llewellyn’s expression, would impose further conditions on these candidate arguments, Gardner (1987), pg.9. For example, it should be shown that each proposition of an argument is *backed* by adequate legal authority or evidence, as in Toulmin’s theory of argumentation, Toulmin (1958), pg.103.

4 A Theory of Issues and Argument Moves

Let me now begin to present the theory of issues, argument moves and clear cases in detail. A set of propositions containing the various interpretations of primary legal sources, to be called Γ ,

⁴In a recent article, Levesque takes a similar view, but orders theories by their *simplicity* rather than their size, where simplicity is defined in terms of the number of unique literals in the theories Levesque (1989).

plays a central role in this theory. This set is expected to be inconsistent, to allow for alternative interpretations, but not every sentence may be a member of this set. Rather, the set is restricted to only “technically correct” interpretations of the primary sources. The theory, at the moment at least, does not further elaborate this notion of technical correctness. It is simply assumed that each proposition in the set of existing interpretations can be shown to be a reasonable interpretation of some primary legal source, such as a statute or case decision. This link to a primary source is similar to Toulmin’s idea of the “backing” for a “warrant”, Toulmin (1958). Or, to use Susskind’s terminology, the set of interpretations is restricted to law “statements” which can reasonably be asserted to be interpretations of particular authoritative law “formulations”, Susskind (1987), pp. 36-37.

This assumption, that a set of alternative reasonable interpretations can be identified, is considerably weaker than legal positivism’s claim that a secondary rule of recognition can be applied to the primary legal sources to find a single, consistent, albeit incomplete, theory of the legal domain of interest.

Notice that the set of interpretations, Γ , is a proper subset of the set of all sentences of some language, \mathcal{L} , as most sentences in \mathcal{L} cannot be backed with references to authoritative legal sources.

Of course, the interpretations included in some finite Γ cannot be considered to be exhaustive of all potentially relevant interpretations of the law. Moreover, if we consider the language \mathcal{L} to consist of all sentences recursively generated from some finite set of predicate and function symbols, as is usual, then the set of all possible interpretations is neither a subset nor superset of any particular \mathcal{L} . It is not a superset, for the same reason that Γ is not equivalent to \mathcal{L} : for any given \mathcal{L} , most sentences in \mathcal{L} cannot be backed with legal authority. But, unlike Γ , the set of all possible interpretations is not a subset of \mathcal{L} , as we want to allow for the possibility of extending the set of predicate and function symbols during interpretation. In natural language, words do not have a fixed meaning, but acquire new shades of meaning during use. In artificial languages such as \mathcal{L} , a unique predicate or function symbol is required for each meaning. Thus, no bounded set of symbols can be adequate for all potential interpretations.

Although a complete account of Fiedler’s constructive view of legal reasoning must surely explain creative legal interpretation, there is as yet no computational theory of legal reasoning which adequately shows how arguments can be constructed from an existing set of competing interpretations. Although a theory of creative interpretation would be ambitious and interesting, it is usually a sensible research strategy to focus on easier questions first.

Thus, the theory here is restricted in just this way: it is only intended to account for reasoning with finite subsets of the complete space of possible interpretations. However, the theory is incremental, in that it prescribes how to modify existing arguments or search for new arguments when new interpretations are asserted; but it does not suggest how to discover these new interpretations.

4.1 Legal Issues

Intuitively, an *issue* is a proposition which is *known* to be *relevant* to the outcome of the case. That is, issues are propositions which are known to make a difference. Questions about facts which can have no bearing on the outcome, however interesting they may be for other reasons, are not issues. Such statements are irrelevant, or simply *non-issues*. Propositions which would make a difference, if certain logical consequences of some existing interpretation of the law were known, but do not make a difference given the inferences which have already been made, are *potential issues*. Finally, because of the open-texture of legal concepts, for example, it may happen that propositions become issues during the process of interpreting the law. We might want to call these *latent issues*.⁵

Thus, to develop a more precise notion of issue, we must also consider what it means for a proposition to be relevant, and what it means to know this. Let us start by hypothesizing the

⁵Potential and latent issues are, of course, issues, in a generic sense. But when the term *issue* is used here, without further qualification, a narrower sense is intended, meaning something like *raised issue*.

existence of a set of sentences, Σ , expressing all possible interpretations of the law. Now, let Γ be a finite subset of Σ . As a main goal of this theory is to account for reasoning with multiple, inconsistent interpretations, let us assume that Γ is inconsistent. It may contain, for example, alternative interpretations of some statute.

Issues are raised only in the process of deciding particular cases. So, let us introduce the concept of a *context*, denoted by Θ , to mean a set of propositions for the interpretations of the law and facts being presently assumed for the sake of argument. The propositions of a context will be sometimes referred to as *assumptions*. I will also speak of *goals* and, of course *issues*. Both are just ordinary propositions, distinguished only by the role they play in argument. As we will see, issues are defined relative to a goal in some context, so it is somewhat misleading to speak generally of the issues of a case. Every case may involve any number of contexts and goals, depending on the current focus of attention. Finally, the term *argument* will be used here to mean a set of propositions which have been shown to entail some goal.

The definitions which follow will depend on an entailment relation \models . However, there are a number of varieties of entailment in the literature. The theory here does not require us to make a specific choice among them. Rather, we only require that the underlying consequence operation satisfy the usual Tarskian properties of inclusion, idempotence and monotonicity, as discussed in, e.g., Alchourron and Martino (1985).

We also require a notion of consistency. Let \perp denote absurdity or contradiction. If $\Psi \models \perp$ then $\Psi \models \phi$ for every proposition ϕ in the language. That is, every sentence is entailed by an inconsistent set of propositions.

Given this terminology, let's now take a first stab at trying to define "issue":

A proposition ψ is an issue, with respect to a goal proposition ϕ , if and only if there exists a set of propositions Δ such that:

1. $\Delta \subseteq \Gamma$,
2. $\Delta \cup \Theta \models \phi$ and $\Delta \cup \Theta \not\models \perp$,
3. There does *not* exist a set Δ' such that $\Delta' \subset \Delta$ and $\Delta' \cup \Theta \models \phi$, and
4. $\psi \in \Delta \setminus \Theta$

In English, a proposition ψ is an issue if and only if we can construct a "theory", Δ , from the current set of known interpretations of the law, Γ , which, together with the assumptions of the current context, Θ , consistently entails the goal, ϕ . Furthermore, Δ must be *minimal*; i.e. there is no subset Δ' of Δ which also, together the assumed propositions, entails the goal. This minimality condition is what insures that ψ really "makes a difference" and is not just some arbitrary proposition added to the argument. The final condition simply prevents us from making an issue out something we have assumed in the current context. Notice that the process of constructing Δ from Γ is a use of abduction, as defined previously.

This may seem to be an intuitively plausible definition of an issue, but there is a problem: it suggests that all the minimal $\Delta \cup \Theta$ entailing ϕ are known. Recall that, intuitively, the issues are those propositions which are known to make a difference with respect to some goal in the context currently being considered. Given a suitably expressive logic, such as first-order predicate logic, entailment is at best only semi-decidable. In such a case, it would not be very realistic to expect an agent to know all the logical consequences of some context.⁶

What to do? My solution is to define "issue" in terms of an entailment relation representing the current state of the argument. This relation, let us call it \models_d , where the subscript d is intended to be suggest "dialog" or "debate", must have the following properties:

⁶Similar considerations, incidentally, lead Hector Levesque to distinguish between explicit and implicit belief in, for example, Levesque (1989). Theories of belief describe the subjective beliefs of a single idealized, rational agent. Here, however, our concern is prescribing what parties to an argument should be expected to know about the logical consequences of various sets of assumptions at each stage of a debate.

1. If $\Delta \models_d \phi$ then $\Delta \models \phi$; and
2. For any ϕ , the set of all minimal Δ such that $\Delta \models_d \phi$ and $\Delta \not\models_d \perp$ is tractably decidable.⁷

That is, the theory asserts that, at each stage of an argument, the minimal, consistent arguments *from which a proof of some goal has already been constructed* can be efficiently computed. These are reasonable requirements for a concept of argument state. The previously made arguments should be efficiently retrievable from memory. The only task which may require substantial effort is computation of the minimal, consistent subsets of these arguments. However, these minimal arguments should also be able to be cached, to be updated only as needed when further arguments have been made.

To avoid potential confusion here: \models_d is *not* a particular relation which has been proved to satisfy these properties. Rather, it is an *hypothesis* that argument states should (since this is a normative theory) indeed satisfy these properties, and a *specification* which any computational model for the theory must be shown to fulfill. In section 5 below, one way to implement \models_d is described, using an ATMS reason maintenance system.

The revised definition of issue simply replaces \models with \models_d . For the sake of completeness, here it is:

Definition 1 (Issue) *A proposition ψ is an issue with respect to a goal proposition ϕ in a context Θ if and only if there exists a theory Δ such that*

1. $\Delta \subseteq \Gamma$,
2. $\Delta \cup \Theta \models_d \phi$ and $\Delta \cup \Theta \not\models_d \perp$,
3. There does not exist a theory Δ' such that $\Delta' \subset \Delta$ and $\Delta' \cup \Theta \models_d \phi$, and
4. $\psi \in \Delta \setminus \Theta$

Let us use the previous version to define *potential issues*:

Definition 2 (Potential Issue) *A proposition ψ is a potential issue with respect to a goal proposition ϕ in a context Θ if and only if there exists a theory Δ such that*

1. $\Delta \subseteq \Gamma$,
2. $\Delta \cup \Theta \models \phi$ and $\Delta \cup \Theta \not\models \perp$,
3. There does not exist a theory Δ' such that $\Delta' \subset \Delta$ and $\Delta' \cup \Theta \models \phi$,
4. $\psi \in \Delta \setminus \Theta$, and
5. ψ is not an issue

Alternatively, it would have been possible to call the potential issues simply “issues”, as in our draft definition, and refer to what we are now calling issues as “raised issues”, or something similar. This work was motivated by Anne Gardner’s AI program for spotting issues in law school examination questions. In that context, student performance is evaluated in terms of how many issues have been spotted. So, given my current definition of issue, a student might want to complain that she or he has found all the issues, as only those potential issues which have been spotted *are* issues! But let’s not start picking nits here; in this theory, students spot potential issues.

To summarize, an issue is a relative concept depending on: 1) the current state of the argument, 2) the current context, which includes stipulated facts as well as uncontested interpretations of the law, and 3) a particular goal proposition.

Let us complete this section by also defining *latent issues*.

⁷Actually, I would like to require \models_d to be *practically* decidable, as tractability alone, meaning that an algorithm of polynomial time complexity exists, is no guarantee that the algorithm will actually be useful given the time constraints of the application.

Definition 3 (Latent Issue) *A proposition ψ is a latent issue with respect to a goal ϕ in a context Θ if and only if there exists a theory Δ such that*

1. $\Delta \subset \Sigma$ but $\Delta \not\subset \Gamma$, and
2. $\Delta \cup \Theta \models \phi$ and $\Delta \cup \Theta \not\models \perp$,
3. There does not exist a theory Δ' such that $\Delta' \subset \Delta$ and $\Delta' \cup \Theta \models \phi$,
4. $\psi \in \Delta \setminus \Theta$, and
5. ψ is not a potential issue

Latent issues differ from potential issues only in two respects: 1) The theory Δ abduced to entail the goal proposition ϕ includes at least one new interpretation from Σ and 2) ψ is not a potential issue, i.e. cannot be made an issue by using only existing interpretations in Γ . Notice that ψ can be a latent issue even though ψ itself was part of Γ , as the complete argument using ψ to show ϕ may depend on other propositions in the new interpretation.

Latent issues, as a matter of principle, cannot be computed; as soon as a new interpretation is found and can be talked about, it becomes part of Γ . Thus, as soon as we have created some new interpretation from which an argument using ψ showing ϕ *can* be made, ψ becomes at least a potential issue. (Once the argument actually *has* been made, ψ is an issue, not a potential issue.)

4.2 Argument Moves

As issues have been defined with respect to some particular state of an ongoing argument, it might be interesting to try to extend this approach to account for certain kinds of argumentation *moves*. In his work on case-based legal reasoning, Kevin Ashley distinguishes between three principal *roles* for precedents in legal argument: “cited cases”, “distinguished cases” and “counterexamples” Ashley (1989). Similarly, I would like to distinguish “arguments”, “rebuttals” and “counterarguments” here. Rebuttals and counterarguments were defined in Gordon (1989), but again I will be redefining them here using the more abstract framework of abduction, rather than defining them directly in terms of the ATMS reason maintenance system.

The term “argument” is used in everyday speech both to mean ongoing debate or discussion and to mean a proof that a particular set of statements entails some desired conclusion. Although I will continue to use the term in both senses, the definition which follows is intended to capture only the latter use: a set of statements *which has been demonstrated* to entail some goal proposition. This definition is, of course, not terribly interesting. Its only purpose is to provide a basis for defining rebuttals and counterarguments.

Definition 4 (Argument) *A theory Ψ is an argument for some proposition ϕ if and only if $\Psi \models_d \phi$.*

Again, \models_d is used to represent the decidable entailment relation capturing the inference steps which have already been made. If $\Psi \models_d \phi$ then not only does Ψ entail ϕ but a proof of this already has been found and can be reproduced at will.

Intuitively, a *rebuttal* dispels the force of an argument. If an argument is brought forth showing ϕ , the purpose of the rebuttal is to show that ϕ is nonetheless not the case. At first glance there may seem to be a number of possible approaches for achieving this. One could show, for example, that the proponent’s reasoning was faulty, i.e. that $\Psi \not\models_d \phi$.⁸ However, in this case the proponent has failed to present an argument, given the above definition. One could modify the definition to distinguish between valid and invalid arguments, but I propose instead to restrict the notion of a rebuttal to moves which dissipate the force of logically sound arguments, as follows:

⁸I will speak of the *proponent* and *opponent* of some argument, rather than the *plaintiff* and *defendant*, as either party can play both roles at different stages of the argument. The plaintiff, for example, can try to rebut a proposition essential to the defendant’s rebuttal of his initial argument.

Definition 5 (Rebuttal) A theory Δ is a rebuttal to an argument Ψ , if and only if $\Delta \cup \Psi \models_d \perp$.

That is, a rebuttal is a set of additional propositions which have been shown to be inconsistent with the argument being rebutted. If Δ is empty, then Ψ was already inconsistent. Recall that we did not require the proponent to show that his argument is consistent. Everything, including ϕ , is entailed by an inconsistent set of premises, so an inconsistent Ψ would be an argument given the above definition. The burden of showing inconsistency has been placed on the opponent. I think this allocation is fair. Proving inconsistency is intractably hard in general, even in the case of an expressively weak logic such as the propositional calculus. If $\Psi \models_d \perp$ just after an argument Ψ has made, then the argument is obviously inconsistent and it is not asking too much of the opponent to point this out. If inconsistency is not obvious, however, possibly requiring arbitrarily deep reasoning, then requiring the proponent to show consistency would be a severe impediment on the ability to rationally justify claims.

It may seem that we have forgotten about the defeasibility of legal reasoning, Gordon (1988). Shouldn't it be possible to rebut an argument by bringing in additional information, without necessarily showing that this information is inconsistent with the facts and law presented by the proponent? In a nonmonotonic logic, it is possible that $\Delta \cup \Psi \not\models \phi$ even though $\Psi \models \phi$. Recall, however, that we have required \models to be monotonic. As it turns out, this is not a limitation. Conveniently, defeasible reasoning can itself be viewed as abduction using an ordinary, monotonic logic, Poole (1988).

One small blemish here is that rebuttals do not raise issues. Recall that, according to Definition 1, only theories not known to be inconsistent with the current context may raise issues. However, the very purpose of a rebuttal, in this formalization, is to construct an inconsistent argument.⁹ Perhaps the definition of issues should be relaxed in the case of rebuttals.

There is one more kind of argument move to define, *counterarguments*. Intuitively, a counterargument is a special kind of rebuttal. Not only is a theory asserted which contradicts the proponent's argument, this opposing theory entails the opposite conclusion. That is:

Definition 6 (Counterargument) A theory Δ is a counterargument to Ψ if and only if $\Delta \models_d \neg\phi$.

Obviously, as the argument Ψ is known to entail ϕ , $\Delta \cup \Psi \models_d \perp$. That is, any counterargument is also a rebuttal.

Notice that whether some theory Δ is an argument, rebuttal or counterargument depends simply on its role or use. That is, the same Δ conceivably could be used to serve all three functions at various stages of the argument.

4.3 Clear Cases

The sources of uncertainty in legal reasoning, each of which can prevent a case from being a *clear case*, are usually enumerated in about this manner, Susskind (1987), pg. 240, Berman and Hafner (1987):

Vagueness or Open Texture. Vague or unclear terms may have been used. When is an action "reasonable" or a cause "proximate"? This is related to the problem of the "open texture" of terms, first introduced to legal philosophy by Hart (1961). Is a "motor home" or "house boat" a "home" for Fourth Amendment search and seizure purposes, Rissland (1989), or a "bicycle" a "vehicle", to mention Hart's famous example, for the purpose of a regulation prohibiting vehicles from a park. Although related, vagueness and open texture are, strictly speaking, separate problems, Susskind (1987), pg. 187.

⁹In Gordon (1989), I did not include this consistency requirement in the definition of issues, precisely to avoid this problem. However, I now feel that too many issues are raised if consistency is not required. Intuitively, it does not seem plausible to consider a proposition to be an issue when it is *known* to be inconsistent with the statements believed in the current context.

Implied Exceptions. Legal rules may be subject to implicit exceptions. The often repeated example in the literature is *Riggs v. Palmer*, 115 N.Y. 506 (1889), where it was decided that a murderer shall not inherit from his victim, despite the explicit rules controlling the devolution of property.

Normative Gaps. An interpretation of the law may be indecisive in some case. This can occur when a concept used in the law is open-textured. If the law refers to homes, and the case involves a house boat, then the law, literally interpreted, is inconclusive, as it is unclear whether the meaning of home includes house boats. Contrast this with the park vehicles example, where there is no question (in my mind at least) that the context independent, literal meaning of vehicle includes bicycles.

Competing Statutes or Precedents. There may be any number of relevant, conflicting precedents. Berman and Hafner cite *Marsh v. Alabama*, 326 U.S. 501 (1946) and some of its successors, which concern when there is a First Amendment right to demonstrate on privately-owned property, Berman and Hafner (1987), pg. 25. Statutes may also conflict, for example when it is unclear which state's laws are applicable, Berman and Hafner (1987), pg. 28.

Unclear *Ratio Decidendi*. The *ratio decidendi* of a case does not declare itself and cannot be uniquely determined. Any case, and therefore its *ratio*, can be described at different levels of abstraction. Courts may later broaden or narrow the *ratio* of a precedent to permit or prevent its application to the case at hand.

Structural Ambiguity. The logical structure of rules formulated in natural language is often unclear. Many examples of this sort are documented in Allen and Saxon (1987).

Examining this list, we can see that the sources of uncertainty in legal reasoning can be subsumed into three more abstract categories: underdeterminedness, overdeterminedness, and problems of interpretation.

4.3.1 Underdeterminedness

When no decision follows deductively from an interpretation of the law then it is *incomplete* or *underdetermined*. That is, suppose the goal proposition is ϕ and the interpretation being considered is Ψ , where Ψ is a consistent subset of Γ . Then Ψ is underdetermined with respect to a context Θ if and only if $\Psi \cup \Theta$ entails neither ϕ nor $\neg\phi$. That is, it is underdetermined just when $\Psi \cup \Theta \not\models \phi$ and $\Psi \cup \Theta \not\models \neg\phi$.

We may also speak of the underdeterminedness of multiple, competing interpretations of the law in some context. The set of known interpretations, Γ , is underdetermined in a context just when no *consistent* subset of $\Gamma \cup \Theta$ entails either ϕ or its opposite:

Definition 7 (Underdeterminedness) *A set of interpretations Γ is underdetermined with respect to a goal ϕ in a context Θ if and only if, for all consistent subsets Δ of Γ :*

1. $\Delta \cup \Theta \not\models \phi$ and
2. $\Delta \cup \Theta \not\models \neg\phi$.

The consistency of Δ is required here, unlike in the definition of an issue, as otherwise Γ would never be underdetermined in the sense intended here. As Γ is expected to be inconsistent, due to competing precedents, e.g., and all sentences are entailed by an inconsistent set of sentences, we would need only let Δ equal Γ to show that Γ is not underdetermined with respect to any goal.

Underdeterminedness can result from open-textured or vague concepts, the failure of the legislature to anticipate the particular case, or the expression of the *ratio decidendi* of a precedent at a too concrete level of abstraction. Implied exceptions do not result in underdeterminedness. Without the exception, an interpretation is decisive, but in the opposite direction.

4.3.2 Overdeterminedness

Whereas underdeterminedness leaves us undecided for want of arguments supporting either outcome, *overdeterminedness* creates uncertainty by providing multiple defensible arguments: either outcome can be defensibly justified. Again, the complete set of known interpretations Γ is expected to be inconsistent. However this general inconsistency of the known interpretations of the law does not necessarily create uncertainty with respect to some goal in a particular context. We say that Γ is overdetermined *with respect to some goal* just when there are *consistent* subsets of Γ which, together with the propositions of the context, entail opposite outcomes:

Definition 8 (Overdeterminedness) *A set of interpretations Γ is overdetermined with respect to a goal ϕ in a context Θ if and only if there exist consistent subsets of Γ , let's call them Δ and Δ' , such that:*

1. $\Delta \cup \Theta \models \phi$, and
2. $\Delta' \cup \Theta \models \neg\phi$.

Notice that Γ cannot be both overdetermined and underdetermined with respect to the same goal in some context.

Referring again to the original list of sources of uncertainty in legal reasoning, above, overdeterminedness can be caused by, for example, multiple inconsistent precedents, or by multiple plausible interpretations of the logical structure of some statute.

4.3.3 Interpretation

Overdeterminedness and underdeterminedness are *logical* sources of uncertainty. In both cases, uncertainty is a property of a set of interpretations of the law and facts, Γ , relative to some goal in a particular context. However, the largest source of uncertainty in legal reasoning is the possibility of *new* interpretations. For example, a new implicit exception may be found, or a precedent may be reinterpreted, creating a new *ratio*.

New interpretations, however, are not simply an additional source of uncertainty. Methods of interpretation, such as those described in Dickerson (1975), are primarily tools for *resolving* existing uncertainty. Where the law is underdetermined or overdetermined with respect to some issue, further creative interpretation can resolve the conflict.

Let us characterize the uncertainty of interpretation as follows. Recall that Σ is the set of all possible interpretations. As such, Σ can never be completely known. Again, let Γ be a finite set of interpretations. Then, new interpretations in Σ create uncertainty with respect to an issue ϕ just when Γ is neither underdetermined nor overdetermined with respect to ϕ (i.e. a unique result is determined by Γ) and a new interpretation exists such that the union of the new interpretation and Γ is either overdetermined with respect to ϕ or determines a result which is contrary to the previous result.

To help define this formally, let us first introduce two utility functions. Let $\mathbf{d}(\Gamma)$ be the unique decision determined by Γ for ϕ in Θ . If Γ is either underdetermined or overdetermined with respect to ϕ in Θ , then \mathbf{d} is undefined. And let $\mathbf{c}(\psi)$ be the *complement* of ψ . That is, if $\psi = \phi$ then $\mathbf{c}(\psi) = \neg\phi$ and if $\psi = \neg\phi$ then $\mathbf{c}(\psi) = \phi$.

Definition 9 (Uncertainty due to Interpretation) *The set of all possible interpretations Σ is a source of uncertainty with respect to some goal ϕ in a context Θ , given some set of interpretations Γ , if and only if:*

1. Γ is neither overdetermined nor underdetermined with respect to ϕ in Θ , and
2. There exists a ψ in $\Sigma \setminus \Gamma$ such that:
 - (a) $\{\psi\} \cup \Gamma$ is overdetermined with respect to ϕ in Θ , or

$$(b) \mathbf{d}(\Gamma) = \mathbf{c}(\mathbf{d}(\{\psi\} \cup \Gamma))$$

The main point of this definition is that if Γ was either underdetermined or overdetermined with respect to the issue, then new interpretations in Σ can only resolve uncertainty, not create it. A further point is that Σ can create uncertainty even though the decision after including the new interpretation is determined, so long as the new decision is the complement of the prior one. This is, of course, a common pattern in legal reasoning. A lower court may apply a rule using, for example, a literal interpretation, just to be overruled by an upper court which finds, e.g., an implicit exception. This is what is referred to as a “retrospectively hard” case in Susskind (1987), pg. 248. Such cases are clear both before and after the new interpretation, but with opposite results.

Let me now attempt to define the concept of a clear case. The basic idea is that a case is clear just when the main issue is determined by the known interpretations of the law or if the only issues which prevent a definite decision are issues of fact.

Definition 10 (Clear Case) *A case is clear with respect to a goal ϕ , given a set of interpretations of the law Γ , a set \mathcal{F} of competing interpretations of the facts of the case, and a context Θ of uncontested facts, if and only if:*

1. *The set $\Gamma \cup \mathcal{F}$ is neither overdetermined nor underdetermined with respect to ϕ in Θ , or*
2. *The set of potential issues is a subset of \mathcal{F} .*

This definition uses the previous definition of a potential issue. Thus this is an *objective* definition of clear cases. At any particular stage of an ongoing argument, some, if indeed not most, of these potential issues may not yet have been identified. Also, the uncertainty caused by latent issues in Σ does not play a role here. Rather, clearness is defined relative to a identified, finite set of interpretations. Thus, “retrospectively hard” cases are also clear cases. Classifying a case as “clear” does not imply that its ultimate result is completely certain. The possibility of a new interpretation causing a contradictory result remains. Admitting this possibility does not render the distinction between clear and hard cases useless or meaningless.

Why is the distinction between clear and hard cases of such importance? There appear to be at least two reasons. The law, as a normative system, is expected and intended to guide conduct. This is only possible to the extent that persons are able to predict the legal consequences of their actions. If there could be no clear cases, then there would be serious cause to doubt whether the law can serve this normative function. The possibility of retrospectively hard cases does not, by itself, diminish this function of the law. Planning is always subject to risks; the risk that the law will be reinterpreted in some way detrimental to the interests of an agent is just one among many. All that planning requires is that not all risks be equal. It is sensible to assume that one’s own case will be covered by the body of previously decided cases, and probably not raise interesting new issues requiring creative interpretation.

Clear cases are also important for limiting the power, or “discretion” if you prefer, of judges. A decision contrary to the result determined by existing interpretations in a clear case is, or should be, subject to an especially high level of justificatory explanation.

The theory of clear cases presented here preserves both of these purposes of the concept. In fact, by supporting reasoning with multiple interpretations of the law, the theory diminishes the risk of planning errors. If, for example, the precedents diverge as to whether a house boat is a home for Fourth Amendment search and seizure purposes, a police officer may decide to avoid the risk of having evidence excluded by first obtaining a warrant. In constructing a consistent theory of search and seizure law, according to the principles of legal positivism, one would have to either leave the question completely open, causing excessive uncertainty, or choose between the precedents, giving the impression of certainty where there is none.

5 The Computational Model

The computational model for our theory of legal reasoning includes, among other things, a theorem prover for full first-order predicate logic, a programming language for expressing control knowledge and a reason maintenance system for managing dependencies between contexts. In this paper, I would like to focus on the ATMS reason maintenance system, as it is of critical importance when reasoning with multiple interpretations of the law. In the next two subsections, I first describe the ATMS and then show how it is used to implement the theory of issues and argument moves.

5.1 The ATMS Reason Maintenance System

The principal job of a reason maintenance system is to cache inferences which have been made. At the cost of memory, a reason maintenance system increases efficiency, trading space for time. When faced with a theorem to prove, the problem solver can first check whether the theorem has already been proved.

One of the first problem solvers to include a reason maintenance system was AMORD, de Kleer et al. (1977). Two members of the AMORD team were Johann de Kleer and John Doyle. Doyle went on to write his doctoral thesis on the problem of truth maintenance, Doyle (1979). Doyle's system, although it was quite significant, will not be of further concern to us here. Rather, we will focus our attention on the *Assumption-based Truth Maintenance System* (ATMS) designed by his partner on the AMORD team, de Kleer (1986).¹⁰

First, let me introduce some of de Kleer's terminology:

- The ATMS manages dependencies between data of some arbitrary type. The only restriction is that the type must include a distinguished element, denoted by \perp .
- Each datum is associated with a unique ATMS *node*.
- Each node is designated to be either an *assumption*, a *premise*, or a *derived node*.
- Inferences are recorded by asserting *justifications*. These may be written as $\Psi \Rightarrow \varphi$, where Ψ is a set of nodes and φ is a node. The left-hand side of a justification is called its *antecedent*, the right-hand side its *consequent*.
- A set of assumptions is called an *environment*.
- The primary task of the ATMS is to determine whether a node *holds* in an environment, given the current set of justifications. Let us use $\Delta \vdash_{atms} \phi$ to mean that the node ϕ currently holds in the environment Δ .
- An environment in which the node for \perp holds is a *nogood*.

Because of the rather unconventional terminology, it may be unclear how an ATMS can be used to cache inferences made by some logical calculus. Let's recast the ATMS in logical terms. A node is a proposition. Assumptions, premises and derived nodes are just propositions with somewhat different roles, as will be discussed shortly. \perp is the proposition denoting absurdity or falsity, as usual. A justification, $\Delta \Rightarrow \varphi$, is a derived inference rule stating that φ is derivable from Δ . A *context* is a set of propositions derivable from the union of the premises and some environment, i.e. some subset of the assumptions. That is, the premises are in every context. The \vdash_{atms} relation is defined by a set of justifications. As each justification is a derived inference rule, \vdash_{atms} is a subset of the derivability relation \vdash of the underlying calculus, which must be monotonic.

¹⁰De Kleer was not the first to design and implement an assumption-based reason maintenance system, this honor belongs to Martins and Shapiro (1988); but as fate will have it, de Kleer's system has been the most influential, perhaps because of his description of an efficient implementation.

De Kleer’s major contribution is an efficient implementation allowing the \vdash_{atms} relation to be defined incrementally by asserting justifications. Determining whether $\Delta \vdash_{atms} \phi$ is true given the current set of justifications is tractably decidable.

Now we can better see how an ATMS can be used to increase the efficiency of a theorem prover. Every time the prover makes a new inference, $\Gamma \vdash \varphi$, it is recorded by asserting a justification, $\Gamma \Rightarrow \varphi$, thus extending \vdash_{atms} . Later, when trying to show that some $\Gamma \vdash \varphi$ holds, the prover first consults the ATMS to see if $\Gamma \vdash_{atms} \varphi$ currently holds. Because \vdash_{atms} is tractably decidable, this is a useful tactic, assuming the general derivability relation \vdash of the theorem prover is only semi-decidable, as is usually the case.

The ATMS is useful not only for improving the efficiency of deduction, but also for abduction. Given a set of justifications, the ATMS computes, for each node ϕ , the set of all minimal, consistent environments from which ϕ is known to be derivable. This set of environments is called the *label* of the node.

To see how this works, let’s describe the implementation of the ATMS in a bit more detail. De Kleer’s uses terms such as “nodes” and “labels” to describe the ATMS, because it is implemented as a directed network. There is a node in the network for each assumption and each justification. For each justification $\Psi \Rightarrow \phi$ there is a link in the network from each assumption in the antecedent Ψ to the node for the justification, and a link from the justification to the consequent ϕ . Each node’s label is, roughly speaking, a set of environments (i.e. a set of a set of assumptions). Each time a justification is asserted, the network is extended and the labels are updated. For each node ϕ , the updating algorithm guarantees the following four properties of its label:

Correctness. For each Γ in the label, $\Gamma \vdash_{atms} \varphi$ holds.

Consistency. For each Γ in the label, $\Gamma \not\vdash_{atms} \perp$.

Completeness. If $\Delta \vdash_{atms} \varphi$ then there exists a Γ in the label such that $\Gamma \subseteq \Delta$.

Minimality. It is not the case that $\Gamma_1 \subseteq \Gamma_2$, for any two environments Γ_1 and Γ_2 in the label.

It should now seem plausible that $\Delta \vdash_{atms} \phi$ is indeed tractably decidable. Because of the monotonicity of \vdash_{atms} , it is only necessary to check whether there is a Γ in the label of ϕ such that $\Gamma \subseteq \Delta$.

Recall that the task of abduction is to find the smallest, consistent subsets of some set of propositions Γ entailing a proposition of interest ϕ . If the entailment relation of interest is completely defined by \vdash_{atms} , then these consistent subsets are just the environments of the label of ϕ ! Thus, in this special case at least, abduction is also tractably decidable: the problem is reduced to simply looking up the label of ϕ in a table.¹¹

5.2 Issue Spotting and Argument Moves with the ATMS

Now we are in a position to describe how the ATMS can be used to spot issues. Recall that issues were defined in terms of a hypothetical inference relation, denoted \models_d , representing the current state of the ongoing argument. \models_d was specified as an inference relation having two requirements: it must be a subset of the entailment relation of the underlying logic, and the set of all minimal sets of propositions from which a proposition is derivable must be tractably decidable. As discussed above, \vdash_{atms} is the just such a relation. Thus, \models_d is implemented using \vdash_{atms} in the computational model.

Using \vdash_{atms} , a proposition ψ is an *issue* raised by a goal proposition ϕ in a context Θ , assuming that all ATMS assumptions are propositions from Γ , the set of interpretations, if and only if there exists an environment Δ in the label of ϕ such that $\psi \in \Delta \setminus \Theta$.

¹¹ In Levesque (1989) it is suggested that computing the labels in the first place from a set of justifications is exponentially complex, in the worst case, but I am concerned here with merely looking up the labels after they have been incrementally updated.

The ATMS also allows the theory of argument moves to be efficiently implemented. A set of assumptions Ψ is an *argument* for some proposition ϕ just when Ψ is a superset of some environment in the label of ϕ . Similarly, Δ is a *rebuttal* to an argument Ψ if and only if $\Delta \cup \Psi \cup \Theta$ is in the label of \perp ; that is, if $\Delta \cup \Psi \cup \Theta$ is a nogood. Recall that an argument need not be shown to be consistent with Θ , the stipulated context. If it was known to be inconsistent, then $\Psi \cup \Theta$ was a nogood, and there is a trivial rebuttal in which Δ is empty. Finally, Δ is a *counterargument* to Ψ if and only if Δ is a superset of an environment in the label of $\neg\phi$.

The procedures just described do not *generate* arguments, rebuttals or counterarguments. Rather, they *recognize* them after they have been made. To generate arguments from the set of interpretations Γ , the computational model searches heuristically through the space of subsets of Γ , looking for Δ such that $\Delta \vdash \phi$. The ATMS is also of great assistance here. As we are only interested in arguments which are not known to be inconsistent (to avoid trivial rebuttals) we need never again consider any argument which is a superset of any of the nogood environments in the label of \perp . Large portions of the large search space of arguments can be efficiently pruned. Finally, when we switch the focus of our attention from one argument to another, as is required by most heuristic search strategies, such as best-first search, the caching service performed by the ATMS permits all the inferences made while previously exploring the consequences of some subtheory of the current argument to be inherited without further computation.

The rest of the terms defined in the theory, such as the concept of a clear case, have yet to be implemented in the computational model. However, it seems safe to assume that the problem of recognizing a clear case, even from a finite set of interpretations, will be intractably hard, or even only semi-decidable, depending on the underlying logic chosen. For example, if the logic chosen is full classical first-order logic, in which deduction is only semi-decidable, then one cannot expect abduction to be decidable, let alone tractable.

One of the terms defined in the theory, the concept of a *latent issue* simply cannot be computed, for the reason mentioned just after the definition: once an interpretation has been created, issues which depend on it are no longer latent, but at most *potential*. One might want to make a similar argument for potential issues, but as the \models relation constructed during a dispute represents the memory of those involved and this memory is not necessarily available to others involved in a new dispute, a program that computes all possible issues from a given set of interpretations can also be considered to have computed the potential issues.

6 A Comparison with Gardner's Program

In her book on Artificial Intelligence and legal reasoning, Anne Gardner describes an AI System capable of spotting issues in offer and acceptance law school examination questions, Gardner (1987). Although it is difficult to do her complex program justice in this space, at the risk of oversimplifying I would like to describe briefly its architecture, as I understand it. Then I will compare the system with my theory and computational model of issues and argument moves.

Gardner represented the law of offer and acceptance as well as relevant common sense knowledge, in three ways:

1. An *augmented transition network* is used to represent the relationship between the various legal events of offer and acceptance law, such as offer, counteroffer, acceptance, and so on. As the net is traversed while analyzing the examination question, the current node is labeled with a propositional representation of the facts of the current event in the question, using the MRS logic programming system, Genesereth (1987). The arcs represent alternative possible legal interpretations of such facts.
2. A goal literal is associated with each arc of the transition network. An arc can be traversed only if the goal is satisfied by the facts in the label of the node. The predicate of the goal literal is defined with MRS "rules", that is with sets of logical formulas, as are other legal and common sense concepts. Such rules can be organized into sets and marked as complimentary

or competing. The rules in competing sets represent alternative interpretations of some primary legal source.¹²

3. The system also includes a limited facility for representing and reasoning with cases. Each of the stored cases is considered to be a clear, prototypical case for some predicate. The cases are used “when the rules run out” to decide whether a goal raises a hard question. If the case matches exactly the facts of the current case, then the question is clear, otherwise it is hard. Cases are also used to realize defeasible reasoning. If a case contradicts a general rule, the question is considered hard and both solutions are recorded in the program’s output.

It is not necessary to try to describe here in much detail how the program works. Its output is a two-tiered decision tree. The top-level tree is a tree of *contexts*, where a context is a set of propositions. Each node in the top-level decision tree includes its own lower-level decision tree for determining which branch to follow in the top-level tree. Each of the nodes in these trees represents a legal or factual *issue* spotted by the program. The easy questions have all been deductively decided during the generation of these trees; they are not reflected in the output.

Alternative theories of the case can be recovered from the top-level tree. The context of each node in the tree represents one interpretation of the case up until the event to be interpreted at that point in the tree. The context associated with a node, it seems, inherits and extends the context of the node’s parent. Thus, the contexts of the leaf nodes contain complete alternative theories of the case.

Thus, the output of the program contains not only the issues spotted in the problem, but also arguable alternative resolutions of these issues.

Now I would like to discuss one way to reimplement Gardner’s program using my computational model of the theory of issues and argument moves. We start by supposing that the relationship between the relevant legal events of offer and acceptance law, such as offers and counteroffers, has been represented as a first-order theory, rather than using an ATN. There are a variety of ways to accomplish this. For example, one could directly translate the ATN into a set of predicate logic sentences. Nodes could be represented by terms, arcs by a binary predicate. The preconditions of an arc could be represented as a set of Horn clauses.

“Common sense” knowledge and competing interpretations of the legal sources would be represented using Poole’s abductive approach to defeasible reasoning, Poole (1988), as extended in Brewka (1989) to allow multiple levels of defaults. The effect of Gardner’s limited use of cases can be achieved with propositions representing defaults using Poole’s approach. Each case becomes a single proposition representing the *ratio decidendi* of the case. Both this approach and Gardner’s offer the same possibilities and suffer from the same limitations; they do not support reasoning by analogy. Gardner chose to view a conflict between a case and some rule as raising a hard question. This can be achieved in our representation by giving the defaults representing cases the same priority as other defaults. If the cases should override the general rules, they could be given higher priority.

Supposing that the law of offer and acceptance has been represented along these lines, in Γ , potential issues in an examination question could be spotted using the following procedure:

1. Represent the facts in the examination question as Θ , a consistent set of uncontested propositions.
2. Let the goal proposition ϕ be $\exists x \dots \text{contract}(x)$.
3. While there is time left and theories can be found:
 - (a) Search for an argument Δ , a subset of Γ , such that $\Delta \vdash \phi$, using the ATMS to avoid multiple searches of parts of the space found to be inconsistent.

¹²This is where Gardner departs from legal positivism, or at least the pure form of positivism we having been using in this paper.

- (b) Cache this argument as a justification in \vdash_{atms} .
- (c) If an argument has been found, switch roles to that of the opponent, and try to find a rebuttal or a counterargument.

4. Inspect the labels in ϕ , $\neg\phi$ and \perp in \vdash_{atms} to retrieve the issues found.

In this procedure, the student (or program) alternates roles between the plaintiff and the defendant, but there may be other strategies, such as first trying to find a number of alternative arguments for one party before switching roles.

Gardner's program not only spots issues but, as mentioned, it generates a decision tree for "deciding" the case. (This is quotation marks, because I do not want to suggest that one should actually decide a case in this manner. After all, only known interpretations of the law have been considered in generating the decision tree.) It would not be difficult to devise a comparable procedure using the information contained in \vdash_{atms} . One could select an issue, decide it, add the choice made to the context Θ , and then recompute the issues, without considering further arguments in Γ . These steps would be repeated until there were no more issues.

The alternative theories of the case are represented in the leaf nodes of the decision tree in Gardner's system. In our approach, these theories are contained in the label of the goal proposition. An important point is that the ATMS guarantees that these alternative theories are minimal. No proposition in a theory is not relevant for deciding the goal. It is unclear whether Gardner can make this claim for her program.

It is also unclear whether our system would spot the same issues as Gardner's program or, from the opposite perspective, whether her program can be viewed as a correct implementation of my theory of issues. It may be that her program identifies too many issues; that some of the "issues" do not really make a difference to the ultimate goal of the case. I suspect that her program may determine issues only locally, with respect to subgoals in some rule during backward chaining, without being able to recognize that the particular path chosen is not relevant globally.

7 Conclusion

What has been accomplished? A formal theory of issues has been presented which provides a *post facto* specification of Gardner's issue spotting program. Of course it is not entirely clear whether her program correctly implements the theory. It is probably not possible to decide this question definitely from the rather abstract description of the program in Gardner (1987). However, this problem points out the value of specifications such as this. The theory, I hope, describes quite clearly just what my computational model should compute.

Also, an alternative computational model of the theory has been described, based on de Kleer's ATMS, but it remains to be proven that it correctly implements our theory of issues.

Although the theory is principally concerned with identifying or spotting issues, it can also concern *issue resolution*, at least to the extent this is possible from a fixed set of interpretations. To see this, suppose the current context, Θ , is $\{\beta\}$ and the label of the goal ϕ contains the argument (i.e. environment) $\{\gamma, \beta\}$. Let's call this argument Ψ . Thus, γ is an issue. Now, suppose there is a proposition ρ in Γ stating $\neg\gamma \vee \neg\beta$. It would then be possible to construct a rebuttal, where Δ is $\{\rho\}$, because $\Delta \cup \Psi$ is inconsistent. This rebuttal is capable of resolving the issue. That is, once the rebuttal has been made, if it is found persuasive and Θ is extended to include ρ , then γ is no longer an issue. Thus, one method of resolving issues is to search for rebuttals of this kind.

Certainly some important types of legal reasoning, such as reasoning by analogy, do not appear to be accounted for in this theory of issues and argument moves. More generally, the theory and its implementation are limited to reasoning with a set of alternative interpretations of the law, but does not suggest how interesting new interpretations can be found. However, this is a limitation shared by Gardner's program and, to a much greater extent, legal reasoning systems restricted to a single, consistent interpretation of the legal domain, regardless of whether they are based on mechanical jurisprudence or legal positivism.

The theory could be extended in a number of ways:

- The theory currently assumes that interpretations in Γ are properly backed, but does not itself include a theory of backing. The problem is to find adequate objective criteria for distinguishing between plausible and ludicrous interpretations. Also, the notion of backing needs to be extended beyond mere pedigree or authority to account for completely new interpretations.
- A case should be clear, even when Γ is overdetermined with respect to some issue, if established secondary rules can be applied to choose, as a matter of law, one of the alternative interpretations. For example, rules enacted by a higher authority are preferred to those from a lesser authority, rules enacted later in time are preferred to earlier rules, and more specific rules are to be preferred to more general rules.
- More generally, following Dworkin, it may be that various *principles* nonetheless converge to a definite result, requiring one of the competing alternative interpretations to be selected. Such a case should be a clear case in the theory. Note that this use of principals is distinct from their use to reach a decision contrary to the expected result in a clear case. Similarly, arguments about the purpose of a law should be useful for ordering interpretations.

The computational model is being implemented in Standard ML, a typed functional programming language, Harper et al. (1988), using the Isabelle generic theorem prover, Paulson (1989). We are also designing an application for assembling legal documents in the field of German divorce law, which will use the system.

References

- Alchourron, C. E. and Bulygin, E. (1989). Limits of logic and legal reasoning. In Martino, A. A., editor, *Logica, Informatica, Diritto*, Florence. Consiglio Nazionale delle Ricerche.
- Alchourron, C. E. and Martino, A. E. (1985). Logic without truth. In *Proceedings of the Conference on Expert Systems in Law*, Bologna.
- Allen, L. E. and Saxon, C. S. (1987). Some problems in designing expert systems to aid legal reasoning. In *The First International Conference on Artificial Intelligence and Law*, pages 94–103.
- Ashley, K. D. (1989). Toward a computational theory of arguing with precedents: Accommodating multiple interpretations of cases. In *The Second International Conference on Artificial Intelligence and Law*, pages 93–102. ACM.
- Berman, D. H. and Hafner, C. D. (1987). Indeterminacy: A challenge to logic-based models of legal reasoning. *Yearbook of Law, Computers and Technology*, 3:1–35.
- Brewka, G. (1989). Preferred subtheories: An extended logical framework for default reasoning. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 1043–1048. IJCAI, Morgan Kaufmann.
- Charniak, E. and McDermott, D. (1985). *Introduction to Artificial Intelligence*. World Student Series. Addison-Wesley, Reading, Massachusetts.
- Dickerson, R. (1975). *The Interpretation and Application of Statutes*. Little, Brown and Company.
- Doyle, J. (1979). A truth maintenance system. *Artificial Intelligence*, 12:231–272.
- Dworkin, R. M. (1977a). Is law a system of rules? In Dworkin, R. M., editor, *The Philosophy of Law*, pages 38–65. Oxford University Press.

- Dworkin, R. M., editor (1977b). *The Philosophy of Law*. Oxford Readings in Philosophy. Oxford University.
- Edwards, P., editor (1972). *The Encyclopedia of Philosophy*, volume 1. Macmillan Pub. Co, Inc. & The Free Press.
- Fiedler, H. (1985). Expert systems as a tool for drafting legal decisions. In Martino, A. A. and Natali, F. S., editors, *Logica, Informatica, Diritto*, pages 265–274, Florence. Consiglio Nazionale delle Ricerche.
- Gabbay, D. and Guenther, F., editors (1983). *Handbook of Philosophical Logic, Volume I: Elements of Classical Logic*, volume 164 of *Synthese Library*. D. Reidel.
- Gardner, A. v. d. L. (1987). *An Artificial Intelligence Approach to Legal Reasoning*. Artificial Intelligence and Legal Reasoning. MIT Press.
- Genesereth, M. R. (1987). An overview of metalevel architecture. In *Proceedings of the National Conference on Artificial Intelligence*, pages 119–123, Los Altos. AAAI, Morgan Kaufmann.
- Genesereth, M. R. and Nilsson, N. J. (1987). *Logical Foundations of Artificial Intelligence*. Morgan Kaufmann, Los Altos.
- Gordon, T. F. (1988). The importance of nonmonotonicity for legal reasoning. In H. Fiedler, F. Haft, R. T., editor, *Expert Systems in Law; Impacts on Legal Theory and Computer Law*, pages 111–126. Attempto Verlag, Tübingen.
- Gordon, T. F. (1989). Issue spotting in a system for searching interpretation spaces. In *Proceedings of the Second International Conference on Artificial Intelligence and Law*, pages 157–164. ACM.
- Harper, R., Milner, R., and Tofte, M. (1988). The definition of Standard ML, version 2. Technical report, Laboratory for Foundations of Computer Science, Department of Computer Science, University of Edinburgh.
- Hart, H. (1961). *The Concept of Law*. Oxford University Press.
- Hayes-Roth, B. and Hayes-Roth, F. (1979). A cognitive model of planning. *Cognitive Science*, 3:275–310.
- de Kleer, J. (1986). An assumption-based tms. *Artificial Intelligence*, 28.
- de Kleer, J., Doyle, J., Steele, G. L., and Sussman, G. J. (1977). Explicit control of reasoning. Memo 427, MIT AI Laboratory.
- Kneale, W. and Kneale, M. (1962). *The Development of Logic*. Clarendon Press.
- Leith, P. (1985). Clear rules and legal expert systems. In Martino, A. A. and Natali, F. S., editors, *Logica, Informatica, Diritto*, Florence. Consiglio Nazionale delle Ricerche.
- Levesque, H. J. (1989). A knowledge-level account of abduction. In *IJCAI-89*, pages 1061–1067, Detroit.
- MacCormick, N. (1978). *Legal Reasoning and Legal Theory*. Clarendon Law Series. Oxford University Press.
- Martins, J. P. and Shapiro, S. C. (1988). A model for belief revision. *Artificial Intelligence*, 35(1):25–80.
- Paulson, L. C. (1989). The foundation of a generic theorem prover. *Journal of Automated Reasoning*, 5:363–397.

- Poole, D. (1988). A logical framework for default reasoning. *Artificial Intelligence*, 36(1).
- Rissland, E. (1988). AI and legal reasoning. *AI Magazine*, 9(3):45–55.
- Rissland, E. L. (1989). Dimension-based analysis of hypotheticals from supreme court oral argument. In *The Second International Conference on Artificial Intelligence and Law*, pages 111–120. ACM.
- Susskind, R. E. (1987). *Expert Systems in Law*. Oxford.
- Toulmin, S. E. (1958). *The Uses of Argument*. Cambridge University Press.