

A Design for Reasoning with Policies, Precedents, and Rationales

Ronald P. Loui
Department of Computer Science
Washington University St. Louis, MO 63130
loui@ai.wustl.edu 314-935-6102

Jeff Norman
Seyfarth, Shaw, Geraldson & Fairweather
55 E. Monroe Chicago, IL 60603

Jon Olson
MIT Lincoln Laboratory
244 Wood Street Lexington, MA 02173-9108

Andrew Merrill
Swarthmore College
Swarthmore, PA 19081

1. Relevance of Formal Analysis of Argument.

The studies of Tom Gordon [Gordon87, Gordon89, Gordon93a] and Henry Prakken [Prakken91, Prakken93a, Prakken93b], together with the work of Rissland and Ashley (*e.g.*, [Ashley89a], [Ashley&Aleven91], [Rissland&Skalak91]) show an interest in models of reasoning that directly address the concepts of argument, defeat among arguments, and dialectical processes through which arguments provide warrant for conclusions. This interest is distinct from the longstanding interest, among legal scholars, in models of reasoning based on relevance, deontic, and intuitionistic logics.

Meanwhile, formal work on non-demonstrative argument and defeasible reasoning has led to results that permit systems to be built that perform general-purpose reasoning of this kind, *e.g.*, [Pollock87], [Baker&Ginsberg89], [Simari&Loui92], [Loui *et al.*93].

Three questions arise. First, does the formal work provide any new insights? Second, do the general purpose reasoners provide similar support of legal argument to that provided by the HYPO and CABARET forebears that were designed specifically for the legal domain? Third, does attention to the legal domain provide new ideas for general, formal analysis of argument?

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

This paper reports on a modest attempt to address those questions. First, it sketches a formal model of argument that has emerged over the past decade from work on non-monotonic reasoning in AI and on defeasible reasoning in philosophy. Then it makes observations about fitting the formal model to Rissland-Ashley purposes. Finally, it reports on the design and partial implementation of LMNOP, a program to reason with policies and precedents that is based on the formal model.

2. A Formal Model of Argument.

Our model of argument is one among a number of competitors, about which there is little agreement; however, the number decreases considerably when attention is restricted to those models that provide a **syntactic criterion** for deciding that one argument is better than another, *e.g.*, is based on a more specific rule or a more on-point case. A syntactic criterion is important: in its absence, all reasoning about when an argument suffices as a response to another (thus, about which advocate must make the next dialectical move) will rely on externally supplied information about argument strength. Surely some, or even most, of these determinations will refer to additional information; however, we can be equally certain that programs that rely entirely on this information unduly burden the knowledge representer. HYPO's underlying logic for distinguishing cases and preferring most on-point precedents is an example of a purely syntactic criterion.

Gordon bases his work on the criterion advanced in [Gefner&Pearl92]. Prakken first based his work on [Poole85]. The following model is the last alternative, which mixes ideas from [Poole85], [Loui87], and [Pollock87]. This model is also the only one that explicitly seeks to model policy arguments and adjudications thereof.

The following is intended to be brief but precise. The later discussion only requires that the general ideas be understood.

We suppose a language L , and a meta-language ML . For present purposes, it suffices to suppose that L is propositional, but that propositions have a predicate-term syntax. We can restrict ourselves to 1-place predicates (what Ashley calls *factors*) and to terms that are names of cases, including *cfs*, the current fact situation. We also suppose the relation >— in ML , which relates two sentences in L , $p \text{>—} q$, when one sentence, p , is *reason* for another sentence q . That is, arguments for p can be extended into arguments for q by citing this reason relation.

A sentence in L might be “*is-faculty(a)*” and a sentence in ML might be

“*is-faculty(a)*” >— “*may-attend-faculty-meetings(a)*”.

Quantification is possible with Quine (single) quotation,

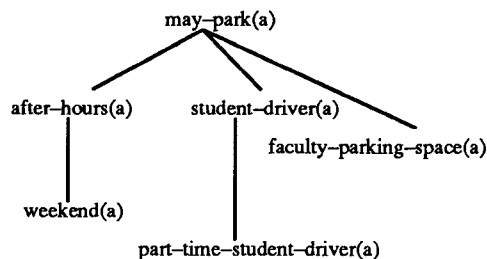
for all $\tau \in \text{Terms}_L$: ‘*is-faculty(\tau)*’ >— ‘*may-attend-faculty-meetings(\tau)*’.

Note that the language is not first-order since quantification does not occur in L and instead is performed in ML . Knowledge of a case is the conjunction of propositions whose predicate-term form refers to the name of the case:

“*is-grad-student(b) & teaches-grad-class(b) & may-attend-faculty-meetings(b)*”.

An **argument** is a pair $A = \langle T, h \rangle$ where h is the **conclusion** of the argument and T is a set of pairs $\langle p, q \rangle$ such that $p \text{>—} q$. Arguments must be grounded in **evidence** (indefeasible knowledge), $K \subseteq L$, which is the set of sentences about which there is no dispute. To be **grounded**, there must be a consistent defeasible derivation from K to h using all and only T . Such a **derivation** is a sequence of sentences in L that either introduce sentences in K , use deductive inference on preceding sentences, or introduce q when p is derivable, and $\langle p, q \rangle \in T$. Arguments are actually best rendered as trees, such as the argument that uses the rules “*weekend(a)*” >— “*after-hours(a)*”, “*part-time-student-driver(a)*” >— “*student-driver(a)*”, and “*after-hours(a) & student-driver(a) & faculty-parking-space(a)*” >— “*may-park(a)*”, grounded in the evidence “*weekend(a) & part-time-student-driver(a) & faculty-parking-space(a)*”.


Two arguments **disagree** when their conclusions are deductively inconsistent. $\langle T_1, h_1 \rangle$ disagrees with $\langle T_2, \text{not}(h) \rangle$. An argument A_1 **counters** another argument A_2 when A_2 has a subargument with which A_1 disagrees. An argument that counters the argument above could have as its conclusion any of “*not(faculty-parking-space(a))*” or “*not(student-driver(a))*” or “*not(after-hours(a))*”, or even the disjunction of all three.




An argument A_1 **defeats** another argument A_2 iff A_1 is more specific than every subargument of A_2 with which A_1 disagrees. According to [Poole85], $A_1 = \langle T_1, h_1 \rangle$ is more **specific** than $A_2 = \langle T_2, h_2 \rangle$ when there are conditions under which A_2 can be used, but not A_1 , and not vice versa; *i.e.*, when there is an activator of A_2 -but-not- A_1 , and there is no activator of A_1 -but-not- A_2 . An **activator** of A_2 -but-not- A_1 (an asymmetric activator) is a sentence in L that could have been evidence (whether it was in fact evidence or not), with which a consistent defeasible derivation of h_2 exists using $T_1 \cup T_2$, where the derivation is non-trivial, and with which the same does not exist for h_1 , *mutatis mutandis*. The condition for **non-triviality** is given in a larger report [Loui *et al.*93]. That report also proves that specificity can be checked quite easily: instead of searching all of L for activators and non-activators, identify the top rules of each argument: the rules that are not used to derive another rule’s antecedent. Conjoin the antecedents of these top rules: usually there is just one top rule. If there is an asymmetric activator, then this conjunction will be an asymmetric activator. Note that this is slightly more complex than saying that specificity can be determined by comparing the antecedents of the top rules.

Two consequences of this conception of defeat are that **more specific rules** are preferred to less specific rules, and **more direct arguments** are preferred to less direct arguments, *ceteris paribus*.

An uncountered argument has the power to **interfere** with another argument and to **support** its conclusion. This is the base case. Two arguments that disagree without either being more specific also interfere with each other, when there are no additional counterarguments to either. An argument,



The left argument defeats the right argument.
 Activators are enclosed in solid ellipses.
 There is an asymmetric activator.




The argument on the left is preferred
 because it uses a rule that is
 more specific.



The argument on the left is preferred
 because it is more direct.

all of whose counterarguments lack the power to interfere,
 can also support its conclusion. [Simari&Loui92] gives a
 labeling procedure based on [Pollock87] that makes this in-
 ductive definition effective.



The argument on the left is reinstated
 by the argument on the right, because
 the latter defeats the middle argument.

The left and right arguments are sup-
 porting, while the middle argument is
 defeated.

A resolution-based C-language implementation of this sys-
 tem, for a first-order L , is available from the authors (also
 reported in [Loui et al.93]).

Syntactically-based defeat is not necessarily the only
 source of preference among arguments that disagree. Ex-
 plicit prioritization occurs naturally in corpora of rules, and
 also naturally arises through consideration of jurisdiction.
 If a prioritization of rules is supplied, it can be used with the
 syntactic criterion, with one ordering supplementing the
 other. Work of this kind is pursued by [Simari91], [Vree-
 swijk93], [Baader&Hollunder93], and [Grosz92].

3. Some Issues of Fitting The Model to Legal Reason- ing.

☞ Do precedents suggest rules, or do they merely permit
 analogies?

Insofar as rules are taken to be of the form $p \succ q$, prece-
 dents suggest rules. First, this is because analogies seem to
 defer to rules when there is conflict, while this seems not
 necessarily the case when reasoning from precedent (case)
 conflicts with reasoning from policy (statute). Second, this
 is forced by the criterion of defeat in order to favor more on-
 point analogies without explicit prioritization.

One of the shortcomings of the model of argument just pre-
 sented is that it cannot support certain kinds of analogies,
 which we shall call **1-projections**. The proper analysis of
 1-projections seems to be as follows (we use the language of
 Pollock):

- a. $P(\text{source}) \ \& \ Q(\text{source}) \ \& \ Q \text{ is-1-projectible-from } P$;
- b. therefore (from a), defeasibly, *being a P is defeasible
 reason for being a Q*;
- c. $P(\text{target})$
- d. therefore (from b and c), defeasibly, $Q(\text{target})$.

This kind of analogy, which appears to be the most common
 analogy in common-sense reasoning, can be countered in
 any of the following ways:

(i) providing an argument for $\text{not}(Q(\text{target}))$: usually be-
 cause this is known; or because there is a 1-projection based
 on a different similarity, $R(\text{source}) \ \& \ R(\text{source}_2) \ \& \ \text{not}(Q(\text{source}_2))$; or a more specific similarity, $R(\text{source}) \ \& \ R(\text{source}_3) \ \& \ P(\text{source}_3) \ \& \ \text{not}(Q(\text{source}_3))$;

(ii) providing an argument against the defeasibly concluded
 defeasible rule, that *being a P is defeasible reason for being
 a Q*: usually because the source is not representative, $P(\text{source}_4) \ \& \ \text{not}(Q(\text{source}_4))$; or in statistical cases, be-
 cause there is a statistical argument that *it is not the case that
 being a P is defeasible reason for being a Q*.

When there is a policy, ' $P(\tau)$ ' \succ — ' $\text{not}(Q(\tau))$ ', presumably this constitutes grounds for knowing ' $\text{not}(P(\tau))$ ' \succ — ' $Q(\tau)$ ', which defeats part (b) of the 1–projection. If we claim that arguments from precedent are analogical, then they seem not to be analogies of the 1–projection kind. Since the model cannot support 1–projections, this is good news.

How, then, do we support the preference of more on–point precedents? It cannot be because there are rules of the form:

$"P(s) \ \& \ Q(s) \ \& \ \underline{P(t)}" \ \succ$ — " $Q(t)$ ",

because in the presence of

$"P(s_2) \ \& \ R(s_2) \ \& \ \text{not}(Q(s_2)) \ \& \ \underline{P(t)} \ \& \ \underline{R(t)}" \ \succ$ — " $\text{not}(Q(t))$ ",

neither is more specific, which is not the desired result. The underlined portions of each rule bear the specificity relation, but the rules as wholes do not.

Instead, it is simpler simply to take a case,

$P(c) \ \& \ Q(c) \ \& \ R(c) \ \& \ H(c)$,

to suggest a combinatorial number of rules,

$'P(\tau)' \ \succ$ — ' $H(\tau)$ '

$'Q(\tau)' \ \succ$ — ' $H(\tau)$ '

$'R(\tau)' \ \succ$ — ' $H(\tau)$ '

$'P(\tau) \ \& \ Q(\tau)' \ \succ$ — ' $H(\tau)$ '

and so forth. There should however be limits to the minimally and maximally specific antecedents, and these will be addressed by providing rationales in the next section.

☞ How can HYPO's factorable factors be represented?

A significant assumption of HYPO, which Ashley acknowledges, is that factors can be determined *a priori* to be favorable or unfavorable to an issue. For example, *agreed-not-to-disclose()* is inherently pro–plaintiff, while *secrets-disclosed-to-outsiders()* is inherently pro–defendant, in arguments regarding trade secrets violations. Moreover, interactions are regular: the combination of two favorable factors is always favorable. That is, factors can be treated as if they are independent, or even **factorable**.

Suppose a case, c_0 *pro*, is decided to have property h *pro* (I will add the subscript *pro* or *con* to a predicate when it is a useful reminder in the present discussion), with the facts f_1 *con* and f_2 *pro*; that is, f_1 is inherently a factor suggesting $\text{not}(h)$, while f_2 inherently suggests h . One case in question, c_1 , exhibits f_1 *con*, f_2 *pro*, and f_3 *pro*. It, too, should be decided to have

h *pro*, since the interaction of f_1 *con* and f_2 *pro* (with a *pro* result) can be factored from the contribution of f_3 *pro*. Another case in question, c_2 , exhibits just f_2 *pro* and f_3 *pro*. It should be decided to have h *pro*, because f_2 *pro* and f_3 *pro* are factorable, and f_2 *pro* has one less $\text{not}(h)$ *con* factor compared to c_0 *pro*. This behavior is abstracted in [Clark88] and reviewed in [Loui89].

This assumption is useful in practice, but is not a general feature of analogical, statistical, or defeasible reasoning (where irregular interactions are permitted). How, then, can we support such an independence of factors?

The case c_0 *pro* suggests the rule: ' $f_1(\tau) \ \& \ f_2(\tau)$ ' \succ — ' $h(\tau)$ '. To say that f_3 is inherently *pro*– h is to say that any rule in which the consequent is ' $h(\tau)$ ' can be enriched to be a more specific rule, ' $f_1(\tau) \ \& \ f_2(\tau) \ \& \ f_3(\tau)$ ' \succ — ' $h(\tau)$ '. This is not actually needed for HYPO–like behavior in c_1 since $f_3(c_1)$ can just be ignored: $f_1(c_1) \ \& \ f_2(c_1)$ suffices to use the rule, and apparently no counterarguments are introduced by adding f_3 *pro* (enriching the antecedent would be required though, if for instance, there were a case in which f_1 *con* and f_3 *pro* had been decided *con*).

Moreover, to say that f_1 is inherently *con*– h is to say that any rule in which the consequent is ' $h(\tau)$ ' and f_1 *con* appears in the antecedent can have f_1 *con* dropped from the antecedent. So the rule ' $f_1(\tau) \ \& \ f_2(\tau)$ ' \succ — ' $h(\tau)$ ' is suggested by the case, as is the rule ' $f_2(\tau)$ ' \succ — ' $h(\tau)$ '. This is required in order to produce HYPO–like behavior in c_2 , since f_3 *pro* may not suffice by itself to provide warrant for h *pro*. It is an inherently *pro* factor, but perhaps it is below the minimum knowledge required to decide *pro*. (Note that our example does yield that f_2 *pro* suffices to provide warrant for h *pro*, but this is because of the combination of factors in c_0 , not because we may create such a rule when f_2 is declared to be inherently *pro*.)

This account of HYPO's factors makes Ashley's assumption seem to be a strong one.

☞ Does superior directness apply to chains of arguments from precedents, to chains of policies, and to mixed chains?

It is possible to chain arguments so that the conclusion of one argument is used in the next argument: this is one way of viewing a subargument. Our syntactic criterion for argument preference says that more direct chains are preferred. This preference appears to make sense for arguments from precedents, though there seems to be no general recognition of this preference, such as the recognition that rules are preferred that have more specific antecedents.

At issue is an argument based on two successive uses of precedents, such as:

is-adjunct-faculty(c₀) decided to be *is-faculty(c₀)*;

is-adjunct-faculty(cfs); hence, *is-faculty(cfs)*;

is-faculty(c₁) decided to be *can-attend-faculty-meetings(c₁)*;

hence, *can-attend-faculty-meetings(cfs)*;

compared with a single, more direct argument from precedent:

is-adjunct-faculty(c₂) decided to be *not(can-attend-faculty-meetings(c₂))*;

is-adjunct-faculty(cfs); hence, *not(can-attend-faculty-meetings(cfs))*.

Do we always prefer the latter argument in virtue of its superior form? This is a question best left for domain experts, and perhaps to legal scholars.

It seems certain that without supplement, directness is dubious for policy chains particularly because one rationale for a policy is that it compresses a chain of existing policies. For example, suppose we adopt

'monday6am(τ) & restricted-space(τ) & non-permit-car(τ) >— 'may-park(τ)'

because of the argument based on the following rules is compressed:

'monday6am(τ) >— 'weekend(τ)';

'weekend(τ) & restricted-space(τ) & non-permit-car(τ) >— 'may-park(τ)'.

This is one way, among many, that policies get adopted (a second rationale for policies is also discussed in the next section; we do not exhaust utilitarian grounds for policies, or other grounds).

By producing a new policy on these grounds, an argument that ought to be defeating might not be; namely, an argument based on the policy:

'weekend(τ) & restricted-space(τ) & non-permit-car(τ) >— 'may-park(τ)',

which is more direct than the previous argument. This was called the problem of stability of theories in the context of inheritance: if, under K , p is defeasibly entailed, then does K retain all its defeasible entailments when $K >— p$ is adopted as a new policy? Any logic of defeasible reasoning that allows directness defeat will not be stable.

If policies are adopted because they compress arguments, then there is a problem with directness defeaters, which threatens the ability to determine automatically what kinds of challenges an argument can withstand. The representer of knowledge is asked, under this regimen, to exhibit “canonical parsimony” [Grosz93]; rules should be written so that there is no doubt about whether they could have been elongated into a chain of more primitive policies.

The use of rationales tempers this concern.

First, we permit an argument to be attacked by attacking the rationales of its rules. In the case of compressed rules, this allows attacks on the elongated form of the argument. So arguments can be made artificially more direct *prima facie*, but the right kind of opposition can still be effective.

Second, an analogy is as direct as it can be: when $P(s)$, $Q(s)$, and $P(t)$, defeasibly, $Q(t)$. By stating the grounds of the analogy, namely the putative similarity, P , the rationale is explicit; there is no intermediary property, e.g. $R(s)$ and $R(t)$, which mediates between possessing P and possessing Q . Thus, reasoning from precedent, *qua* analogy, ought to be compatible with directness defeaters. On the alternate view that cases suggest rules (rather than support analogies), again the ability to use rationales is mitigating. When Q is decided in the case because of an argument that builds on the decision that the case also manifests R , then the situation is like compressed rules. The argument from precedent can be attacked by first recalling the argument that was the rationale of the precedential decision, then by attacking that argument (with its intermediary premises exposed) in the context of the current case.

☞ What are appropriate argument moves?

[Rissland&Skalak91] initiate a welcome discussion about argument moves. This is the discussion to which we have also been led: they consider domain-specific moves while we consider the problem of dialectical protocols in the abstract. For CABARET, Rissland and Skalak report several heuristic rules of control, such as:

If a rule's conditions are not met and one wants the rule to succeed, then broaden the rule;

and

If a rule's conditions are not met and one wants the rule to fail, then confirm the miss

(interestingly, these are defeasible rules describing how to use defeasible rules).

In computing the conclusion reached by argument, counterargument, and adjudication, there are many choices of control. For example, suppose arguments are produced by

backward-chaining search, and search is serial (not all targets of search can be pursued concurrently). Then an argument (where $P(a)$ & $S(a)$ is evidence):

$P(a) >— Q(a)$;

$R(a) >— S(a)$;

$Q(a) \& S(a) >— T(a)$;

can lead to any of a number of responses:

i. find an argument for $not(Q(a))$ or $not(S(a))$ or $not(S(a) \& Q(a))$ or $not(T(a))$ or any other such disjunction;

ii. Attack any of the rationales for the rules used in the argument, such as attacking the rationale of $P(a) >— Q(a)$; e.g., force the acknowledgement of an intermediate premise, $P(a) >— U(a)$; $U(a) >— Q(a)$, then attack that premise by finding an argument for $not(Q(a))$.

iii. retreat to a context in which this argument is no longer relevant; e.g., retreat from the use of $not(T(a))$ as a premise of an earlier argument by finding a different premise.

Any search for argument can chain through rules as well as cases. Each rule with useful consequent, and each case with useful decision presents a new choice: whether to backward chain to show that it applies to the current case. Usually when arguing from cases (but also on occasions when arguing from rules) the strength of the similarity can be varied, and this presents more choice. Suppose a case, t , is located such that $F_1(t) \& F_2(t) \& F_3(t) \& P(t) \& not(Q(t))$. To be more on-point, at least one of F_1, F_2, F_3 must be shown of a . The choice of which to pursue, and how many, is open. Further, if there is a property arguably holding of a , and not yet decided of t , arguments can be sought to establish, *post hoc*, that the property also held of t (e.g., argue that $R(t)$).

Certain dialectical protocols of disputation assist in making these choices. If all that is required is a sufficient response, then any of the F_1, F_2, F_3 will do, and search can be stopped with the first that succeeds. The opposing party is faced with the choice of responding. If the opposition's response creates a context in which it is useful to strengthen the precedent-based argument with the remaining F_i , then search can be continued at that point.

Dialectic obviates the need for each side to produce all of its arguments at a single stage, and it does so by indicating what arguments are minimally acceptable at a given stage. Opposition focuses search.

Not all choice is eliminated in this way, though. Under a fixed protocol (dialectical or not) there will still be the question of what is a good strategy for play, and what information could be used to inform strategic choices.

Our studies have led us to formalize and taxonomize various dialectical protocols for disputation, and to work on a computational and game-theoretic analysis of these protocols ([Loui92], [Loui&Chen92]). Gordon's "Pleadings Game" [Gordon93b] and Cavalli-Sforza and Moore's rules of relevant argumentative discourse [Cavalli-Sforza&Moore92] are further examples of interest in the issue of controlling choice and formalizing protocol.

4. Design of Argument Systems for Policies, Precedents, and Rationales.

Our program, LMNOP, is a two-stage investigation of the issues just discussed.

The first stage sought to recreate the basic functionality if HYPO and CABARET based on the abstract, general theory of argument presented above. At issue was what compromises would have to be made, what functionality could not be recreated, and whether the Poole-based syntactic criterion for argument-preference produced any untoward surprises.

The second stage seeks to study questions about protocol and strategy, and to study the adequacy of certain forms of rationales for precedents and rationales for rules. We view this latter stage as experimental dabbling in difficult issues about which we are not yet prepared to theorize formally. A proper analysis of rationales may differ substantially from what is reported here; nevertheless, it will have been informed by the experience.

The first stage was (essentially) completed in LISP, and we are rewriting the system in GAWK for more flexibility (and the ability to schedule processes under UNIX™) in the second stage.

Sparing the reader a detailed account of lessons learned in the first stage (e.g., that the minimality condition on arguments is unimportant under dialectical regimens, or e.g., that non-boolean dimensions present unanticipated analogies), we summarize: a computer program based on the formal, abstract argument need not differ from the earlier designs of Rissland-Ashley.

To us, this provides evidence of a convergence of thinking. It does not mean that our formal basis ought to be adopted by earlier researchers. It might just mean that the kind of argument we had in mind is more like legal argument than, for example, multi-valued logical argument. It does mean that we have discovered nothing in the Rissland-Ashley paradigm that is disagreeable from a more abstract view of argument. It does mean that there are useful mathematical models that more directly address the interesting reasoning with policies and precedents (and perhaps rationales, too)

than the logics (paraconsistent, relevance, intuitionistic, deontic, dynamic) imported by legal scholars in the past.

4a. Basic examples in LMNOP.

The following simple examples show how the syntactic specificity criterion determines preferences among simple arguments. In each example, lower case is used in the input, and upper case in the output. An exclamation following a property indicates that it was a fact of the current case (evidence, undisputed). The symbol “—<” is for rules, and the symbol “—[” is for cases. In the cases, the first field is the decided property, the second is the list of properties of the case, the third is the minimal similarity (lower bound), the fourth is the maximal similarity (upper bound), and the last (always nil in these examples) is the rationale. In the queries, the first argument is the property in dispute, and the second is a list of properties of the current case.

EXAMPLE 1. This example shows the indeterminate result when two competing arguments can be made from cases and there is no basis for choosing between them (case2 has a relevant factor, but excludes it with the lower bound).

```
DECISIONS
case1 ((not h) ( f1 ) ( f1 ) ( f1 ) nil)
case2 (h ( f1 f2 ) ( f2 ) ( f2 ) nil)

>(argue 'h '(f1 f2))

Getting pro argument = ( H F2! )
H --[ F2! (DEC 0 CASE2)
Getting con argument = ( (NOT H) F1! )
(NOT H) --[ F1! (DEC 0 CASE1)
Neither is better.
Attacking con = ( (NOT H) F1! ) With = ( )
con undefeatable.
Attacking pro = ( H F2! )
Getting another con = ( ) Failed.
Neither side wins.
```

EXAMPLE 2. This example shows how the more specific (more on–point) analogy is preferred; it can be made, unlike in the previous example, because the upper bound has been raised for case 2.

```
DECISIONS
case1 ((not h) ( f1 ) ( f1 ) ( f1 ) nil)
case2 (h ( f1 f2 ) ( f2 ) ( f1 f2 ) nil)

>(argue 'h '(f1 f2))

Getting pro argument = ( H F2! F1! )
H --[ F2! F1! (DEC 0 CASE2)
Getting con argument = ( (NOT H) F1! )
(NOT H) --[ F1! (DEC 0 CASE1)
pro is better.
Getting another con = ( )
Failed.
The argument for H :( H F2! F1! )
```

EXAMPLE 3. This example shows a mix of rules and cases. Unlike Rissland’s canonical argument form, however, the rule gets used in order to make an argument from a case. So this is not an example of case–based reasoning used to resolve open–texture. Neither argument is better because even though the argument for con is more direct, the argument for pro has a more specific rule–antecedent. Technically, it is because (f1 f2) activates con’s argument, and (f1) activates pro’s argument.

```
RULES
(f2 rule1 (f1) nil nil)
DECISIONS
case1 ((not h) ( ) ( f1 ) ( f1 ) nil)
case2 (h ( f1 f2 ) ( ) ( f1 f2 ) nil)

>(argue 'h '(f1))

Getting pro argument = ( H F2 F1! )
H --[ F2 F1! (DEC 0 CASE2)
      F2 --< F1! RULE1
Getting con argument = ( (NOT H) F1! )
(NOT H) --[ F1!
Neither is better.
```

EXAMPLE 4. In this example, the attempt to establish the property (h) fails because a subargument (for f5) is defeated. The subargument is defeated because the disagreeing argument has the more specific rule–antecedent, and both rely on the same analogy to establish (f4).

```
RULES
(h rule0 (f5) nil nil)
(f5 rule1 (f1 f4) nil nil)
((not f5) rule2 (f1 f2 f4) nil nil)
DECISIONS
case1 (f4 ( f3 ) ( ) ( f3 ) nil)

>(argue 'h '(f1 f2 f3))

Getting pro argument = ( H F5 F4 F3! F1! )
H --< F5 RULE0
      F5 --< F4 F1! RULE1
          F4 --[ F3! (DEC 0 CASE1)
Getting con argument = ( (NOT F5) F4 F3! F2! F1! )
(NOT F5) --< F4 F2! F1! RULE2
          F4 --[ F3! (DEC 0 CASE1)
con is better.
Getting another pro = ( ) Failed.
```

4b. Data structures for rationales.

Rationales for Cases.

The simplest structure that records the rationale of a case is a pair of **bounds**, a lower bound and an upper bound, on the similarity of the source and target case.

The next level of detail of a rationale records the **main argument** that ultimately was held to warrant the decision.

We record the rules that participated and reconstruct the argument from the rules.

The most detailed rationale is a (perhaps partial) **record of the dispute** that occurred in reaching the decision. We provide for a list of arguments, whether for pro or con, regardless of each argument's ultimate status as undefeated–interfering, undefeted–supporting, or defeated. Each argument is reconstructed from its rules. The relations among arguments (disagreement, countering, defeat) can be reconstructed.

```
case1
facts: a b c d e f g
decision: not-h
lb: a b
ub: a b c d e
rationale:
  not-h: rule1 rule2
  h: rule3 rule4
  not-o: rule5
decision: not-o
lb: c
ub: a b c d e
rationale:
  o: rule4
  not-o: rule5
```

Arguments based on cases are required to be within upper and lower bounds. An argument based on a case can be attacked by recalling the main argument for the decision and attacking it. If a larger record of the disputation can be recalled, a new counterargument will have to survive existing reinstatements of the main argument, to be effective (or else, a new suite of counterarguments must be advanced, to attack the main argument as well as relevant, existing reinstaters).

Rationales for Rules.

Rationales for rules may be represented in three forms. A rule can **compress** an argument (the parking rule regarding 6am, above, is an example). A rule can **specialize** a general principle. For example, “the young should be encouraged” is a policy. Graduate students are young (apparently in the appropriate sense). Reimbursement for presenting papers is encouragement. Hence, “graduate students should be reimbursed for presenting papers”. Formally, specialization and compression are indistinguishable.

```
rule1
decide: b
if: a e f
opt: g h
r-rationale:
  b: rule12 rule10 rule15
```

If a rule has either of these rationales, the data structure provided is a list of the rules used in the argument that was compressed (or worded differently, the general principle and the rules that supplement its more specific application).

An argument using a rule with this rationale can be attacked by first restating the argument in terms of the uncompressed form, then providing counterargument to the uncompressed argument. The uncompressed argument is invariably weaker because it provides more points with which to disagree, and it is more susceptible to defeat from more direct arguments.

Rules also have lists of “irrelevant” (in the sense of [Geffner&Pearl92]) properties that can enrich the antecedent: that is, they can be added in any combination to the “if” part of the policy, and the result is still a policy. Of course, as evidence is added, a policy that was usable remains usable. This is a different concept: a number of policies exist with the same general form, and they can be conveyed succinctly by giving a lower bound and an upper bound (or maximal increment) on the antecedent.

The third form of rationale for a rule is a list of cases (many of which may be hypothetical cases) which together suggest a **regularity** that the rule seeks to encode. We record the relevant facts of the case, making a distinction between hypothetical and actual cases just in case the decision of a particular case could be questioned (this is not likely, if this is a case, the decision of which provided grounds for legislating a generalization!). A rule with consequent *h* is required successfully to separate cases exhibiting the property *h* from those not exhibiting *h*.

```
rule2
decide: b
if: a e f
opt: g h
c-rationale:
  case1: b a e f g
  case2(h): b a e f h
  case3(h): not-b a e
  case4: not-b a f g
```

An argument using a rule with this rationale can be attacked by first proposing a new rule that also successfully distinguishes the recorded cases exhibiting the property from the cases not exhibiting the property; then, by (1) noting that the new rule no longer applies to the current fact situation (or at least, that the argument that it applies has not been given), or by (2) noting that it applies but is not as specific as originally suggested. In the latter situation, the argument is susceptible to attack by counterarguments that would have been considered less specific on the earlier construal of the rule. An argument using a rule with a regularity–among–cases rationale may also be attacked by adding or deleting cases.

For example, the policy “medical doctors on–staff and attending rounds at the campus student health services (SHS)

may park in the handicapped parking zone (HP)” might allegedly summarize the decisions (case1) “a Ph.D. in anthropology visiting the SHS may not park in HP”; (case2) “a medical doctor not attending rounds at the SHS may not park in HP”; (case3) medical personnel without the M.D. late for an appointment at SHS may not park in HP”. To attack the rationale of this rule, add a case, such as (case4) “medical doctors with non-emergency skills, such as psychiatrists even if on-staff and attending rounds at the SHS, may not park in HP”; or delete a case *e.g.*, claim that (case3) is incorrectly decided; or propose a different regularity consistent with the cases: “doctors, excluding Ph.D.’s, may park in HP”.

Interestingly, we have no guidance from the formal model at this point in determining what would be an effective response to such an attack. For example, could opposing parties engage in a series of restatements of a rule? Which restatements are adequate responses to prior restatements? Many of the appropriate responses will include utilitarian arguments. Attacks on this kind of rationale correspond to reviewing the legislative body’s decision to adopt a rule, and may be risky in legal arguments. However, in less formal, less well-developed domains of reasoning, such as parking rules, departmental policy, guidelines for billing or negotiation, and so forth, in many quasi-legal domains, attacking the regularity putatively encoded by the rule might be apropos.

Detailed rationales are hard to find. Computational experience with these structures has been limited.

5. Conclusion.

Our investigations, tied to the formal models of argument that aim beyond the legal domain, lag behind those who aim precisely at legal reasoning. Hence, LMNOP does not actually represent an advance in software (not even as prototype) for AI and Law.

Its main theoretical innovation is a first proposal for representing rationales, and we do not have the authority to suggest their adequacy, or importance, for legal reasoning,

Its main technological advantage is the uniform treatment of argument. This permits statistical argument, decision-theoretic argument, best-explanation argument, and other forms of argument to be mixed with policy and precedent-based argument, without forcing a broader view. Our formal model, by the original intention of its originators, includes all of those forms of argument in its purview.

Insofar as AI and Law, or the scholarly milieu of law in general, is interested in logical accounts of reasoning, the LMNOP experience is significant. With [Prakken93b] and [Gordon93a], we find most useful mathematics among the

emerging theories that treat arguments and their interrelations directly.

The prospect of a technology that can fundamentally change our most important social institutions drives our research on policy argument, as it drives, too, the research of those who first pursued artificial intelligence in law. It would seem a mistake, then, not to pay attention to each other’s relevant work.

Acknowledgements.

The first author’s work was supported by NSF R-9008012. The latter two authors were supported on NSF R-57135A and NSF CDA-9123643, respectively. We acknowledge discussions with T. Gordon, M. Mittelman, and L. Wolff.

Bibliography.

- Ashley, K. “Defining salience in case-based arguments,” Proc. IJCAI 1989a.
- Ashley, K. “Toward a computational theory of arguing with precedents,” Proc. ACM ICAIL, 1989b.
- Ashley, K. and V. Alevan. “Toward an intelligent tutoring system for teaching law students to argue with cases,” Proc. ACM ICAIL, 1991.
- Ashley, K. and V. Alevan. “Generating dialectical examples automatically,” Proc. AAAI 1992.
- Ashley, K. and E. Rissland. “A case-based system for trade secrets law,” Proc. ACM ICAIL, 1987.
- Ashley, K. and E. Rissland. “A case-based approach to modeling legal expertise,” *IEEE Expert*, 1988.
- Baader, F. and B. Hollunder. “How to prefer more specific defaults in terminological default logic,” DFKI (Saarbrücken) RR-92-58 (to appear in Proc. IJCAI 1993).
- Baker, A. and M. Ginsberg. “A theorem prover for prioritized circumscription,” Proc. IJCAI, 1989.
- Cavalli-Sforza, V. and Moore, J. “Collaborating on arguments and explanations,” Proc. AAAI Spring Symposium on Producing Cooperative Explanations, 1992.
- Clark, P. “Representing arguments as background knowledge for the justification of case-based inferences,” AAAI Workshop on Case-Based Reasoning, 1988.
- Geffner, H. and J. Pearl. “Conditional entailment: bridging two approaches to default reasoning,” *Artificial Intelligence* 53, 1992.
- Gordon, T. “OBLOG2: A hybrid knowledge representation system for defeasible reasoning,” Proc. ACM ICAIL, 1987.

- Gordon, T. "Issue spotting in a system for searching interpretation spaces," *Proc. ACM ICAIL*, 1989.
- Gordon, T. Doctoral dissertation, forthcoming, 1993a.
- Gordon, T. "The pleadings game," this volume, 1993b.
- Grosf, B. *Updating and Structure in Nonmonotonic Theories*, doctoral dissertation, Stanford U., 1992.
- Grosf, B. Personal communication, 1993.
- Loui, R. "Defeat among arguments: a system of defeasible inference," *Computational Intelligence 3*, 1987.
- Loui, R. "Analogical reasoning, defeasible reasoning, and the reference class," *Proc. Knowledge Representation and Reasoning*, 1989.
- Loui, R. "Dialectic, computation, and ampliative inference," in *Philosophy and AI*, R. Cummins and J. Pollock, eds., MIT Press, 1991a.
- Loui, R. "Argument and belief: where we stand in the Keynesian tradition," *Minds and Machines 1*, 1991b.
- Loui, R. "Process and policy: resource-bounded non-demonstrative argument," *Wash. U. Comp. Sci. CS-TR-92-43*, 1992.
- Loui, R. "How a formal theory can be normative: implementation versus interpretation," *Journal of Philosophy 90*, 1993.
- Loui, R. and W. Chen. "An Argument Game," Technical Report, *Wash. U. Comp. Sci. TR-92-47*, 1992.
- Loui, R., Norman, J., Merrill, A., Olson, J., Costello, A., and K. Stiefvater. "Computing specificity," *Wash. U. Comp. Sci. TR-93-03*, 1993.
- Pollock, J. "Defeasible reasoning," *Cognitive Science 11*, 1987.
- Poole, D. "On the comparison of theories: preferring the most specific explanation," *Proc. IJCAI*, 1985.
- Prakken, H. "A tool in modelling disagreement in law: preferring the most specific argument," *Proc. ACM ICAIL*, 1991.
- Prakken, H. "An argumentation framework in default logic," to appear in *Annals of Mathematics and AI*, 1993a.
- Prakken, H. *Logical Tools for Modelling Legal Argument*, doctoral dissertation, Free University Amsterdam, 1993b.
- Rescher, Nicholas. *Dialectics*, SUNY Albany, 1977.
- Rissland, E. "AI and legal reasoning," *AI Magazine 9*, 1989.
- Rissland, E. "Artificial intelligence and law," *Yale Law Journal 99*, pp. 1957-1981, 1990.
- Rissland, E. and Skalak, D. "Interpreting statutory predicates," *Proc. ACM ICAIL*, 1989.
- Simari, G. "Plausible reasoning," draft, 1992.
- Simari, G. and R. Loui. "A Mathematical Treatment of Defeasible Reasoning and its Implementation," *Artificial Intelligence 52*, 1992.
- Vreeswijk, G. *Studies in Defeasible Argumentation*, doctoral dissertation, Free University Amsterdam, 1993.