

How to Buy a Porsche: An Approach to Defeasible Decision Making Preliminary Report

Gerhard Brewka and Thomas Gordon

GMD, Schloss Birlinghoven
53757 Sankt Augustin, Germany
{brewka,gordon}@gmd.de

1 Introduction

In the Zeno project at GMD, we are investigating computational models of decision making in groups. In the spirit of the Pleadings Game [6], we are planning to implement a mediating system for decision making procedures. The purpose of the system is to regulate the rights and obligations of each group member in the process of decision making, to identify illegal moves in the language game, and thus to promote fairness of the process.

The Pleadings Game is a mediating system for a particular kind of procedure: identifying the issues of a legal conflict. In contrast, the Zeno system will be a generic tool for implementing a class of mediating systems. To test the generality of Zeno, we will implement three prototype applications: 1) a reconstruction of the Pleadings Game; 2) a system for supporting the negotiation of contracts; and 3) a system for mediating the planning deliberations of a city council. Pleading is a two-party adversarial process, where the goal is to identify the issues of a legal conflict. Contract negotiation is a two-party cooperative process, where the parties have equivalent roles. The city council scenario is an example of group decision making with a more differentiated role structure; e.g. the mayor has more power than the other council members.

We view group problem solving and decision making as regulated theory construction. To design a mediating system, choices have to be made at three layers. The foundation, or bottom layer, is the choice of logic. The logic specifies the notion of theory to be used. It provides a language for expressing theories and a derivability relation specifying whether or not some proposition is in the theory. The next layer defines the moves of the group decision-making procedure. Basic moves include adding and deleting propositions from a theory. In the envisioned applications there are more abstract moves, such as making claims, arguments, and offers. Finally, at the highest level, the rules of the decision-making procedure are stated. This is the level which distributes rights and obligations, such as the burden of proof. Whereas

the second level defines the space of possible moves, the top level defines what is permissible.

In this extended abstract, we will focus on the bottom layer and present a new logic for defeasible qualitative decision making. Our goal is to identify and model the structures one typically uses in argumentation about the pros and cons of alternative positions proposed as solutions to some problem. Our approach rests on the following five basic assumptions:

1. Argumentation is defeasible. There may be arguments for or against a proposition, or for or against a certain choice, but this should not render the whole system inconsistent or irrational.
2. Decision making is governed by preferences among the arguments supporting conflicting propositions. No single principle, such as specificity, is sufficient for resolving all conflicts among arguments. Rather, the appropriate principles are a part of knowledge of the domain.
3. As in the Pleadings Game, preferences themselves should be derivable and subject to debate. That is, preferences are to be represented within the logical language.
4. The arguments pro and con some choice should be weighed against each other. As pointed out, e.g., by [8], it should be possible for two weak arguments together to defeat a counterargument that is stronger than each of them alone.
5. As it would be difficult to assign numerical costs to arguments, facilities should be provided for expressing preferences qualitatively.

The Pleadings Game was built on top of Geffner and Pearl's logic of conditional entailment, [5]. Conditional entailment uses specificity alone to resolve conflicts. Although a method was found to model various other conflict resolution principles in terms of specificity, such as the principle that later rules have priority over earlier ones, it is not possible using conditional entailment to make specificity subordinate to some other principle. There are domains where this is required. In the law, for example, Federal law has priority over State law, even when the State law is more specific. Also, conditional entailment does not provide a convenient way to encode the idea that weak arguments may join together to defeat some stronger one.

To overcome these problems and satisfy the above requirements in a more direct way, we have developed a new logic, which has its roots in Brewka's preferred subtheories [2] and, more precisely, an extension by Gadi Pinkas called penalty logic [9]. Intuitively, a preferred subtheory of a possibly inconsistent set of premises is a maximal consistent subset of the premises that contains as many of the preferred pieces of information as possible. Penalty logic generalizes this idea by associating with each premise an integer representing the cost of not taking the premise into account. The preferred maximal consistent subsets are then those that minimize the cost of the rejected formulas. In the context of constraint satisfaction similar ideas have been described by [7].

Penalty logics are more general than the standard preferred subtheory approach, in particular they offer the possibility that several weak arguments can beat a strong counterargument. Nevertheless they have two drawbacks:

- They require the user to specify exact cost values.
- They do not offer the possibility to reason about preferences.

For these reasons our approach generalizes penalty logics even further, to allow for qualitative representations of the value of applying a default rule. Moreover, we apply ideas from the Pleadings Game and [3] to represent preferences within the logical language.

The rest of this extended abstract is organized as follows: Section 2 defines the new logic, to be called Qualitative Value Logic. Section 3 discusses how to represent choices and make decisions using the logic, using an example about choosing a car. Section 4 gives another example which demonstrates how rules about priorities may be represented. Finally, Section 5 discusses some problems and issues regarding the logic that we will be attending in the Zeno project in the immediate future.

2 Qualitative Value Logic

In this section we first recall some definitions from earlier papers. For the sake of simplicity we restrict ourselves to finite propositional theories. Let F and D be, possibly inconsistent, sets of formulas. The pair (F, D) is called a default theory. Intuitively, F represents the certain knowledge and D is a set of default rules. $F \cup D'$ is a subtheory of (F, D) iff D' is a maximal F -consistent subset of D . Subtheories are exactly what Poole calls scenarios in [10].

Given a preference relation on subtheories we can define a credulous and a skeptical inference relation.

Definition 1 *A formula p is a credulous (skeptical) consequence of (F, D) iff p is classically entailed by at least one (all) of the preferred subtheories of (F, D) .*

Various ways of defining the preference relation on subtheories have been investigated in the literature [2, 9, 1]. The approach which is most relevant for us here is Gadi Pinkas' penalty logic [9]. The idea of penalty logic is to associate with each rule in D an integer representing the cost of violating the default. The user has to specify a cost function from D to the integers. Equivalently, one can instead specify an evaluation function, for the value of applying a default rule. We prefer this latter alternative, to stress that the logic is intended to be used for reasoning about value judgments.

Definition 2 *Let $T = (F, H)$ be a default theory, and v its associated value function. A subtheory (F, H') of T is preferred iff the sum of the values of the elements in H' is maximal.*

Penalty logics require an exact specification of costs or values. It would be more convenient to be able to describe preferences in qualitative terms. Moreover, penalty logics do not allow us to reason about the preferences within the logic. We will therefore extend penalty logic in several ways.

First, we assume that our underlying logic contains axioms for simple integer arithmetic. This includes an axiomatization of the integer operation $+$ and the

predicates = and > with their standard meanings. Since we want to be able to refer to hypotheses, we assume the existence of a naming function that generates a unique term for each hypothesis. The names of hypotheses can be used arbitrarily as terms in default theories. Moreover, we use a reserved function symbol, v , to represent the value of a default.

The definition of preferred subtheories now proceeds in two steps. In the first step an arbitrary subtheory E of (F, D) is generated. During the generation of this subtheory the intended meaning of $>$, v , and the default names is ignored.

The second step then, basically, eliminates E if it cannot be reconstructed in a way that satisfies the priority constraints in E itself. Let us make this intuition precise:

Definition 3 *Let $T = (F, D)$ be a default theory, n its associated naming function, and E a subtheory of T . An evaluation function **val** is compatible with E iff*

$$E' = E \cup \{v(n_i) = x_i \mid d_i \in D, n(d_i) = n_i, \mathbf{val}(d_i) = x_i\}$$

is consistent.

Definition 4 *Let $T = (F, D)$ be a default theory, n its associated naming function, and E a subtheory of T . E is an optimal subtheory of T iff there is a value function v for D such that*

1. v is compatible with E , and
2. E is a preferred subtheory (in the sense of penalty logic) of T , using v as an evaluation function.

Note that our formalism allows for qualitative as well as numerical representations of the importance of pieces of information. Given two rules r_1 and r_2 we can represent qualitative information such as $v(r_1) > v(r_2)$, but we can also require that the value of some rule be a particular numeric value, e.g., $v(r_1) = 142$.

3 Decision Making with Qualitative Value Logic

Let us discuss how to represent and solve decision problems using this logic. Using the terminology of issue-based information systems [4], there are three kinds of information to represent:

Issues These are the questions to be decided or goals to be achieved.

Positions These are the alternative solutions which have been proposed for resolving an issue or achieving a goal.

Arguments These are assertions about the properties or attributes of each position, which speak for or against choosing it.

In our version of this framework, it will also be possible to make assertions about the relative strengths of the arguments.

Issues, positions and arguments will all be modelled by sentences of propositional logic. To make it easier to express the relationship between an issue and its positions, we will define a *choices* macro. $\text{choices}(n, \{p_1, \dots, p_n\})$ means that p_1 to p_n are the only ways to achieve the goal n and that it is not possible to select more than one alternative. Precisely:

$$\begin{aligned} \text{choices}(n, \{p_1, \dots, p_n\}) &\equiv \\ \neg n &\supset (\neg p_1 \wedge \neg p_2 \dots \wedge \neg p_n) \wedge \\ n &\supset (p_1 \wedge \neg p_2 \wedge \dots \wedge \neg p_n) \vee \\ &\quad (\neg p_1 \wedge p_2 \wedge \neg p_3 \wedge \dots \wedge \neg p_n) \vee \\ &\quad \dots \vee \\ &\quad (\neg p_1 \wedge \neg p_2 \wedge \dots \wedge p_n) \end{aligned}$$

The *choices* macro basically allows us to represent the possible alternative solutions in terms of an OR-tree. The extension to AND/OR-trees is straightforward. For OR-trees each leaf represents a possible position. For AND/OR-trees positions are maximal sets of leaves consistent with $\{tree, root\}$, where *tree* is the logical representation of the tree and *root* its root.

We assume that information about the properties of the different options is given as additional facts in a qualitative value logic theory, that is we are interested in optimal subtheories of a theory

$$(\{tree, root\} \cup F, D)$$

Note that optimal subtheories do not necessarily determine unique positions. To ensure this we can add the leaves of the AND/OR tree as hypotheses and give them cost 0. This has the effect that every preferred subtheory contains a position and that positions requiring more components are not rejected for this reason.

Let us consider an example. Suppose we want to buy a car, and the alternatives being considered are an American or European car, or a Toyota. The choice of American cars is between a Chrysler or a Chevrolet. The alternative European cars being considered are a Porsche and a Volvo. In our approach, these alternatives would be represented by the following expressions:

$$\begin{aligned} &\text{choices}(\text{car}, \{\text{American}, \text{European}, \text{Toyota}\}) \\ &\text{choices}(\text{American}, \{\text{Chrysler}, \text{Chevrolet}\}) \\ &\text{choices}(\text{European}, \{\text{Porsche}, \text{Volvo}\}) \end{aligned}$$

This represents the issues and their positions. The arguments for and against each position are expressed by rules from which the properties of each alternative can be derived, if it is chosen. To continue with our example,

$$\begin{aligned} \text{Porsche} &\supset \text{expensive} \wedge \text{fast} \wedge \text{nice} \wedge \text{safe} \\ \text{Volvo} &\supset \text{expensive} \wedge \text{safe} \wedge \neg \text{fast} \end{aligned}$$

Toyota $\supset \neg\text{expensive} \wedge \neg\text{fast} \wedge \neg\text{nice}$
 Chevrolet $\supset \neg\text{expensive} \wedge \neg\text{fast} \wedge \neg\text{nice}$
 Chrysler $\supset \neg\text{fast} \wedge \text{expensive}$

The properties we would like our final choice to possess are expressed as named hypotheses in D . We use the notation $d_i : p$ to associate hypothesis p with the name d_i :

$d_1 : \text{nice}$
 $d_2 : \neg\text{expensive}$
 $d_3 : \text{safe}$
 $d_4 : \text{fast}$

Finally, value judgments are represented by writing equations and inequations about the values of the different goals:

$v(d_2) > v(d_1)$
 $v(d_2) > v(d_4)$
 $v(d_1) + v(d_4) > v(d_2)$

That is, price is more important than being either nice or fast, but a car which is both nice *and* fast will be preferred regardless of price.

By putting the information about the possible choices (i.e, the decision tree and its root) and their properties in F and the properties of the ideal choice in D , Qualitative Value Logic can be used to derive both the optimal choices and their properties. It is not difficult to see that every value function satisfying the value constraints will prefer the Porsche over the Volvo and Chrysler since the Porsche satisfies the same goals and more (note that it follows from our value constraints that all values are greater than 0). Similarly, the Toyota will be preferred over the Chevrolet, because the Chevrolet, it has been assumed, is not nice.

The final decision is determined by the additional preference information we have: the Toyota satisfies d_2 , but the Porsche satisfies d_1 and d_4 , which together have been assumed to outweigh d_2 . For this reason the single optimal subtheory contains Porsche, the car you always wanted to buy.

4 A Legal Example

The example above about choosing a car to buy does not demonstrate the full power of Qualitative Value Logic, in particular how to represent *principles* for ordering competing rules or arguments. This section will show how this can be done, using the legal ordering principle mentioned in the introduction, that Federal law has priority over a conflicting State law, even when the State law is more specific.

We use predicates for expressing information about the relative specificity and source of some rule:

$federal\text{-law}(r_1)$
 $state\text{-law}(r_2)$
 $more\text{-specific}(r_2, r_1)$

Here, r_1 and r_2 are assumed to be the names of two propositions, using the naming convention introduced previously. The example is intended to mean that r_1 is Federal law, r_2 is State law, and that r_2 is more specific than r_1 .

Let us now express the principle that federal law has priority over state law, and that more specific rules have priority over more general rules. As the version of Qualitative Value Logic we have presented here is propositional, we will use propositional schemata to express generic rules and names. Schema variables are denoted here by x and y .

$$\begin{aligned} \text{lex-superior}(x, y) &: \text{federal-law}(x) \wedge \text{state-law}(y) \supset v(x) > v(y) \\ \text{lex-specialis}(x, y) &: \text{more-specific}(x, y) \supset v(x) > v(y) \end{aligned}$$

We further assume that the following fact solves conflicts among the latter two rules:

$$v(\text{lex-superior}(x, y)) > v(\text{lex-specialis}(y, x))$$

Now if federal-law r_1 and state-law r_2 are compared other things (values) being equal the federal law is preferred, even if the state law is more specific.

5 Coming Attractions

Qualitative Value Logic is an important step towards the goals we stated in the introductory section. Nevertheless, a number of important issues remain to be addressed in the near future:

1. Qualitative Value Logic makes it possible to express information about the value of accepting a certain hypothesis in the logical language. Our two examples have shown that this expressiveness can be used to reason about the degree to which a decision satisfies our goals, and on the other hand to reason about the quality of arguments. This can be viewed as a qualitative treatment of utilities and probabilities, respectively. A qualitative decision theory would obviously have to integrate both notions. This cannot be done by simply throwing the different types of information together. For instance, the default information that for a certain type of car the service is normally bad should not be overridden by giving high value to the goal of having good service. To avoid obvious wishful thinking, a more sophisticated treatment is necessary where the default properties of different options are established first and then evaluated against each other.
2. The problem of testing whether a proposition is derivable from a default theory in Qualitative Value Logic is clearly intractable, even when it is limited to propositional logic, as in this paper. This is so, as finding the maximally consistent subset of the defaults requires testing propositional satisfiability, which is known to be an intractable problem. This is a serious problem, as our intention is to use the logic in a mediating system, where the system must decide in some reasonable period of time whether or not some move is permissible. We have some ideas about how to solve this problem, by using discourse norms to distribute the burden of proof among the participants, rather than requiring the mediating system to solve such problems. This approach was also used in the Pleadings Game.

3. In this abstract we concentrated on the basic logical level of our intended system. The final paper will contain more information about the other two layers of the Zeno system. In particular, we intend use the logical formalism presented here not only for representing knowledge of the domain of discourse, but also to reason about permissions and obligations in the discourse game, that is at the third level.

Acknowledgements

Thanks to our colleague Joachim Hertzberg for interesting discussions on the topic of the paper.

References

- [1] S. Benferhat, C. Cayrol, D. Dubois, J. Lang, and H. Prade, Inconsistency Management and Prioritized Syntax-Based Entailment, Proc IJCAI-93, Chambéry, France, 1993
- [2] Gerhard Brewka, Preferred Subtheories - An Extended Logical Framework for Default Reasoning, Proc. IJCAI-89, Detroit, 1989
- [3] Gerhard Brewka, Reasoning About Priorities in Default Logic, to appear in Proc. AAAI, Seattle, 1994
- [4] Jeff Conklin and Michael Begeman, gIBIS: A Tool for Exploratory Policy Discussion. Proceedings of CSCW '88, Portland, 1988.
- [5] Hector Geffner and Judea Pearl, Conditional Entailment: Bridging Two Approaches to Default Reasoning. *Artificial Intelligence*, 53(2-3):209–244, 1992.
- [6] Gordon, Thomas F., The Pleadings Game: An Artificial Intelligence Model of Procedural Justice, dissertation, TU Darmstadt, 1993
- [7] Guesgen, H.W., Hertzberg, J., A Perspective of Constraint-Based Reasoning, Springer Verlag, LNAI 597, 1992
- [8] Hage, Japp, Monological Reason-Based Logic, Proc. ICAIL-93, Amsterdam, 1993.
- [9] Pinkas, Gadi, Propositional Non-Monotonic Reasoning and Inconsistency in Symmetric Neural Networks, Proc. IJCAI-91, Sydney, 1991
- [10] Poole, David, A Logical Framework for Default Reasoning, *Artificial Intelligence* 36:27–47, 1988.