Incorporating Specificity in Extended Logic Programs for Belief Revision

Berilhes Borges Garcia and **Gabriel Pereira Lopes**
Departamento de Informática
Universidade Nova de Lisboa, Portugal
E-mail:{bbg, gpl}@di.fct.unl.pt

Abstract

In this paper a new operator for agent's implicit belief revision is presented. This operator is conceptually based on the following requirement: an agent should find first an explanation for the observations it makes before it tries to revise its beliefs. As a consequence, the proposed belief revision model does not agree with the principle of minimal change, commonly accepted. A translation of an agent's knowledge base, expressed by means of a conditional defaults set for an abductive extended logic program, is also presented. This translation allows us to make explicit the specificity that exists among defaults and also the logic program allows us to determine the possible explanations for a recently observed fact.

Introduction

The assimilation process of new information in an agent's knowledge base has been usually divided into belief revision (Alchourrón, Gärdenfors, & Makinson 1985), (Gärdenfors 1988), (Boutilier 1991) and belief update (Katsuno & Mendelzon 1991). The first describes how a rational agent should change its beliefs when it believes that the world has not changed but its internal state became inconsistent. The update process is concerned with the description of the changing process in an agent's beliefs when it believes that the world has changed. Recently, several researchers (Boutilier 1997), (Friedman 1997) have suggested that both processes may be understood as being opposite extremes of the broad spectrum of belief change strategies that an agent can adopt.

In this paper a new approach to belief revision is introduced. We may say that the slogan adopted by Li and Pereira (Li & Pereira 1996) to describe their possible causes approach suits our proposal too: What is believed is what is explained. However Li and Pereira define a new belief change operator for dynamic domains, where observed changes are due exclusively to actions that might have happened, i.e. they aim to model the belief update process, while we are concerned with the belief revision problem.

Belief revision operators have been governed by the AGM postulates (Gärdenfors 1988), and have been represented by a syntax based approach (Ginsberg 1986) and by semantic based approach (Winslett 1988). Minimality criterion has there a central role. By minimality criterion it is meant that the agent's belief set, after a revision has been made in order to incorporate new information, should differ minimally from the belief set it had before the revision.

In this paper we assume that minimality is a very strong criterion. Evidential reasoning (reasoning triggered by an observation), on the other hand, requires far more flexible criteria where explanation for recently observed facts has a central role.

According to our proposal, when an agent observes a new fact ϕ, it should look for an explanation for that fact, before it ever tries to change its belief base. If it is possible to find an explanation α for the recently observed fact, the agent must change its belief base in order to incorporate $\phi \wedge \alpha$, computing the logic consequences of this new fact and of its explanation. However if the agent is not able to determine an explanation, its beliefs remain unaltered; this means that the agent assumes a skeptical attitude. This scenario is equivalent to the situation where the agent prefers the information in background to the newly observed fact.

This paper is structured in the following way: Next section introduces a motivational example and the formalism used to describe our domain. Section 3 presents how specificity relations can be computed through the syntactic processing of existing clauses in the agent's knowledge base. Section 4 introduces a translation of specificy relations to a extended logic program with explicit negation. In section 5 we discuss how to incorporate new evidence within the agent's knowledge base. In the final section we draw the main conclusions as well as some directions to be followed in future research.

Domain descriptions

Let T be the specific knowledge an agent possesses about a certain context. T is a pair (K_D, E), where K_D represents the background knowledge about the domain the agent has, i.e. the generic knowledge that very unlikely the agent will change and E represents the contingential knowledge, i.e. the knowledge that it is likely to vary from case to case and along the time axis. The pair $T = (K_D, E)$, is also known, as the agent's knowledge base.

The background knowledge K_D will be represented by means of a set of default clauses[1]; of the form: $\alpha_i \rightsquigarrow \beta_i$. Each default, $\alpha_i \rightsquigarrow \beta_i$, is interpreted as a defeasible rule, i.e., "If α_i then normally / typically β_i". β_i is an objective literal (an atom or its explicit negation), defined from a set $\mathcal{L}$ of arbitrary ground literals. α_i is a propositional formula constructed from $\mathcal{L}$ and the logic connectives $\vee$, $\wedge$ and $\neg$. $\rightsquigarrow$ is a meta-connective, meaning normally / typically. $\mathcal{L}$ has also symbol $\perp$, representing logic falsity. Symbol $\models$ represents the relation of logic consequence.

Example 1 *Take agent A, having the following specific knowledge, regarding a certain domain:*

$$K_D = \left\{ \begin{array}{l} d_1 : cs(X) \rightsquigarrow \neg int(X, lp) \\ d_2 : cs(X) \rightsquigarrow \neg int(X, lin) \\ d_3 : int(X, ai) \rightsquigarrow int(X, lp) \\ d_4 : int(X, ai) \rightsquigarrow \neg int(X, lin) \\ d_5 : int(X, ai) \rightsquigarrow cs(X) \\ d_6 : int(X, pr\_cl) \rightsquigarrow int(X, lin) \\ d_7 : int(X, pr\_cl) \rightsquigarrow int(X, ai) \end{array} \right\} \quad (1)$$

Rules d_i represent the following facts: (d_1) A believes that computer science (cs) students are normally neither interested in learning logic programming (lp), (d_2) nor linguistics (lin); (d_3) students interested in artificial intelligence (ai) are normally interested on learning logic programming, (d_4) but typically are not interested in learning linguistics; (d_5) students interested in artificial intelligence are normally students from computer science; (d_6) students interested in doing their final course project on computational linguistics (pr_cl) are typically interested on learning linguistics; (d_7) students interested in doing its course project in computational linguistics are normally interested on artificial intelligence.

Assume now that agent A is informed that b is a computer science student, which can be represented by the contingential knowledge: $E = \{cs(b)\}$.

This said, we can expect that the deductive closure of K_D, where the connective $\rightsquigarrow$ is substituted by the classical material implication $\supset$[2], generates the following belief set:

$$BS_1 = \left\{ \begin{array}{c} cs(b), \neg int(b, lp), \neg int(b, lin), \\ \neg int(b, ai), \neg int(b, pr\_cl) \end{array} \right\} \quad (2)$$

If after a while A gets evidence in favor of the fact that b is interested in studying logic programming, which is represented by the following formula ϕ,

$$\phi = int(b, lp) \quad (3)$$

How should A's beliefs be modified in order to accommodate this new evidence? Both Winslett and Ginsberg approaches lead, in this example, to the same revision and therefore to the same belief set BS_2:

$$BS_2 = \left\{ \begin{array}{c} cs(b), int(b, lp), \neg int(b, lin), \\ \neg int(b, ai), \neg int(b, pr\_cl) \end{array} \right\} \quad (4)$$

However we conjecture that agent A should alter its belief set in order to incorporate both the recently observed fact and its explanation. Therefore, the approach presented in this paper would lead to the belief set BS_3, as $int(b, ai)$ would be an acceptable explanation for being interested in logic programming:

$$BS_3 = \left\{ \begin{array}{c} cs(b), int(b, lp), \neg int(b, lin), \\ int(b, ai), \neg int(b, pr\_cl) \end{array} \right\} \quad (5)$$

We can easily observe that the belief sets (5) and (4) can not be generated by interpreting the background knowledge K_D as being a first-order logic theory. Thus, the agent should not continue to use some of the present clauses in its knowledge base, or otherwise it would be led to an inconsistent belief set.

So, according to our perspective agent's new belief set should incorporate both the logical consequences of the new observation and the logical consequences of its explanation. This being the case, if $BS \cup \phi$ is consistent, the new belief set will be defined as being $Cn(T \cup \phi \cup \alpha)$, where α is an explanation to ϕ and Cn represents the logic consequence operation.

Specificity Ranking Function

Along of this paper we assume that the least specific defaults are the natural candidates to be blocked in any context. So we need to determine which defaults are more specific and for this we use System Z, proposed by Pearl (Pearl 1990). Tolerance is the key concept for partitioning set K_D, into mutually exclusive subsets $K_{D0}, K_{D1}, ..., K_{Dn}$. Two rules belong to a subset K_{Di} are equally specific. If they belong different subsets they have different specificities.

A rule is tolerated by a set of defaults Δ if the antecedent and the consequent of this rule are not in direct conflict with any inference sanctioned by Δ^*, where Δ^* is a set obtained from set Δ by replacing the meta-connective $\rightsquigarrow$ by the classic material implication $\supset$.

Definition 1 (Tolerance) *(Goldszmidt & Pearl 1996) A rule $\alpha \rightsquigarrow \beta$ is tolerated by a defaults set Δ iff $\{\alpha \wedge \beta\} \wedge \{\varphi \supset \psi \mid \varphi \rightsquigarrow \psi \in \Delta\} \not\models \perp$.*

Based on this concept of tolerance Goldszmidt and Pearl define an interactive procedure that allows the generation of partition of the set default rules in the knowledge base of the agent. This procedure can be defined in the following way:

1. Find the set of rules tolerated by K_D, name this set as K_{D0}.

[1] For simplicity and space reasons, in this paper, we will only consider domains described exclusively by default rules.

[2] Temporarily we are abstracting ourselves from the underlying details of the connective use $\rightsquigarrow$ and of the normality concept involved in the clauses in K_D.

2. Find all rules tolerated by $(K_D - K_{D0})$ and name this set K_{D1}.
3. Repeat step 2 until there is no rule in K_D that has not yet been assigned to one of the sets K_{Di}.

This procedure converges and a partition of the rules of K_D will be obtained. This partition has the following property:

Every default belonging to K_{Di} is tolerated by $\bigcup_{j=i}^{n} K_{Dj}$, where n is the number of the equivalence classes defined by tolerance relation. Rule r_i is less specific than a rule r_j iff $r_i \in K_{Dk}$ and $r_j \in K_{Dl}$ and $k < l$.

From this partition a defaults ranking function can be constructed. This function maps each default to a non-negative integer, representing the specificity of the default. This function is obtained by assigning every default the index corresponding to the partition to which it belongs. So, $Z(\alpha \rightsquigarrow \beta)$ is equal to i if the default belongs to the partition K_{Di}, and Z is the name of the ranking function. So that for whichever two defaults d_i and d_j belonging to the set K_D, if d_i is less specific than d_j then $Z(d_i) < Z(d_j)$.

Example 2 *By determining which default is tolerated by the other defaults in the background knowledge K_D of example (1) we obtain the following partition:*

$$\begin{aligned} K_{D0} &= \{d_1,\ d_2\};\ K_{D1} = \{d_3, d_4, d_5\} \quad (6) \\ K_{D2} &= \{d_6, d_7\} \end{aligned}$$

and so $Z(d_1) = Z(d_2) = 0$, $Z(d_3) = Z(d_4) = Z(d_5) = 1$ *e* $Z(d_6) = Z(d_7) = 2$.

Translating in an extended logic program

The basic idea of this section is to propose a translation ζ of the knowledge base $T = (K_D, E)$ to an extended logic program, with two types of negation: explicit negation and negation by default. This program takes into account the specificity relations implicit among existing defaults in the agent's knowledge base, which were determined by System Z. The semantics of this program will be given by the Well Founded Semantics with eXplicit negation (WFSX) (Alferes & Pereira 1996).

An extended logic program P is a set of rules of the following kind:

$$L_0 \leftarrow L_1 \wedge \ldots \wedge L_m \wedge not\ L_{m+1} \wedge \ldots \wedge not\ L_n \quad (7)$$

Where $0 \leq m \leq n$. Each L_i is an objective literal. An objective literal is an atom A or its explicit negation $\neg A$. The symbol *not* represents negation-as-failure and *not* L_i is a default literal. Literals are objective literals or default literals and $\neg\neg A \equiv A$.

If we assume that more specific information has prevalence over least specific, we would be led to conclude that a certain default rule d should have priority over any least specific default rule d', thus inhibiting the use of d'.

So each conditional default $(d : \alpha \rightsquigarrow \beta)$ belonging to a partition K_{D_i} $(i > 0)$ of K_D, will give rise to two rules (δ_d and φ_d) in extended logic programming:

$$\begin{aligned} \delta_d &: \beta \leftarrow \alpha \wedge not\ \neg\beta \wedge not\ ab_i. && (8) \\ \varphi_d &: ab_{i-1} \leftarrow \alpha. && (9) \end{aligned}$$

And for each default $(d : \alpha \rightsquigarrow \beta)$ belonging to a partition K_{D0} a unique rule will be generated:

$$\delta_d : \beta \leftarrow \alpha \wedge not\ \neg\beta \wedge not\ ab_0. \quad (10)$$

The generated set of rules δ_d and φ_d enables us to capture the specificity pattern that is dictated by the conditional interpretation of defaults and reflected in the partition of the knowledge base of the agent.

Definition 2 (ζ-translation) *The ζ-translation of K_D will be equal to $P_D = \bigcup_{i=0}^{n} \{\delta_d \mid d \in K_{Di}\} \cup \bigcup_{i=1}^{n} \{\varphi_d \mid d \in K_{Di}\}$, where n is the number of the equivalence classes defined by the tolerance relation among defaults.*

Example 3 *Consider the background knowledge K_D defined by (1). The ζ-translation of this knowledge is equal to the following program P_D:*

$$\begin{aligned}
\neg int(X,lp) &\leftarrow cs(X) \wedge not\ int(X,lp) \wedge \quad (11) \\
&\qquad not\ ab0. \\
\neg int(X,lin) &\leftarrow ex(X) \wedge not\ int(X,lin) \wedge not\ ab0. \\
ab0 &\leftarrow int(X,ai). \\
int(X,lp) &\leftarrow int(X,ai) \wedge not\ \neg int(X,lp) \wedge \\
&\qquad not\ ab1. \\
\neg int(X,lin) &\leftarrow int(X,ai) \wedge not\ int(X,lin) \wedge \\
&\qquad not\ ab1. \\
cs(X) &\leftarrow int(X,ai) \wedge not\ \neg cs(X) \wedge not\ ab1. \\
ab1 &\leftarrow int(X,pr_cl). \\
int(X,lin) &\leftarrow int(X,pr_cl) \wedge not\ \neg int(X,lin) \wedge \\
&\qquad not\ ab2. \\
int(X,ai) &\leftarrow int(X,pr_cl) \wedge not\ \neg int(X,ai) \wedge \\
&\qquad not\ ab2.
\end{aligned}$$

Example 4 *Given $\langle K_D, \{cs(b), int(b,ai)\}\rangle$. The obtained translation equals $(P_D, \{cs(b), int(b,ai)\})$, this program has a single model, that preserves the specificity relations between defaults, in which the following objective literals are true:*

$$\{cs(b), int(b,ai), int(b,lp), \neg int(b,lin), ab0\} \quad (12)$$

We can observe that (12) partially coincides with (5). The differences are the objective literals $\neg int(b,pr\_cl)$, which appears in (5) due to contrapositive reasoning; and the literal $ab0$, that forces in (12) the specificity relation between the defaults d_3 and d_1.

It should be stressed the role played by the partition during the revision process. As it was said before, when an agent A faces an inconsistency, motivated by the application of some default conditionals, it should abdicate to use some of these defaults, so that it may restore the consistency of his belief set. Usually, in logic programming the onus of determining which defaults should be blocked is under the responsibility of the programmer. However, when the knowledge base grows it may be difficult to stipulate how the various default clauses interact. The use of System Z, and therefore of the default ranking function, allows us to bring to surface the preference relations that exist between the different conditional defaults include in the agent's knowledge base. Thus defining which defaults should be used and which should be blocked during the revision process.

Proposition 1 (Soundness wrt to Specificity) *The Well Founded Model (WFM) of (P_D, E) is sound with relation to specificity of K_D dictated by System Z, i.e. whenever a literal β belongs to WFM then either $\beta \in E$ or there is at least one default $d : \alpha \rightsquigarrow \beta$ such that $\alpha \in WFM$ and there is no default $d' : \alpha' \rightsquigarrow \neg\beta$ more specific than d such that $\alpha' \in WFM$.*

The ζ-translation of the knowledge base $T = (K_D, E)$ to an extended logic program plays two roles in our framework. In first place it provides a efficient method for computing the logic consequences of the agent's known context. Furthermore, it also suggests a mechanism whereby abductive reasoning can be elaborated by agent; how will be shown in the next section.

Assimilation of new evidence

This section introduces our approach for the problem of incorporating new evidence in an agent's beliefs base. According to our proposal, the first step in the assimilation process of new evidence ϕ by agent A, consists in determining the predictive explanations to the newly observed fact. We will say that α is a predictive explanation to an observation ϕ, when the belief in α is sufficiently strong for inducing the belief in ϕ; in other words by believing in α the agent makes a commitment with ϕ.

So, according to our proposal, α is a predictive explanation to ϕ in a context $T = (K_D, E)$ iff there is least one default $d : \lambda \rightsquigarrow \phi$ such that $\lambda \in WFM$ of $T' = (K_D, E \cup \alpha)$ and $\lambda \notin WFM$ of $T = (K_D, E)$; in other words α is a predictive explanation to ϕ if there is one active default $d : \lambda \rightsquigarrow \phi$ in context $T' = (K_D, E \cup \alpha)$ and it is more specific than any other active default $d' : \lambda' \rightsquigarrow \neg\phi$ in $T' = (K_D, E \cup \alpha)$.

Obviously this first step happens only if the new evidence does not yet belong to the set of beliefs the agent has. If the agent already believes in the new evidence, no action will be carried out, and the belief base of the agent remains unaltered.

Therefore we need to determine the abductive explanations for a newly observed fact ϕ. For that purpose we use an abductive logic framework (Eshghi & Kowalski 1989). In the extended logic programming context, an abductive framework P_A may be obtained, according to Kowalski and Eshgi, in the following way:

$$P_A = \langle P_E, Abd, IC \rangle \tag{13}$$

Where $P_E = P_D \cup \{\sigma \leftarrow \mid \forall \sigma \in E\}$, IC denotes the integrity constraints set and is equal to $IC = \{\bot \leftarrow not\ \phi\}$, E represents agent's contingential knowledge, ϕ represents the new observation to be explained. Abd is the abductive literals set and it is equal to the subset of the antecedents set in the domain description K_D ($Abd \subseteq Ant(K_D) \backslash \phi$). That is, for $\phi = int(b, lp)$ and K_D of example (1), $Abd = \{cs(X), int(X, ai), int(X, pr\_cl)\}$.

Δ is an abductive explanation to ϕ iff: $P_A \cup \Delta \models_{WFSX} \phi$ and $P_A \not\models_{WFSX} \Delta$.

In this way the predictive explanations can be computed by evaluating abductive queries to the extended logic programs. It should be stressed that an observation ϕ may have more than one possible explanation in program P_A, and therefore in the belief base T. So, we will refer to the set of possible explanations in T for ϕ by $\|T \wedge \phi\|$.

We can now define a new operator $\diamond$ for the revision of the belief base T motivated by the new fact ϕ. This operator will generate a new belief base set, which will assimilate the new evidence:

$$T^{\diamond}_{\phi} = \begin{cases} \{T_i \mid T_i = (K_D \cup E_i)\} \\ \quad where\ E_i = E \cup \phi \cup \Delta_i \\ \quad to\ each\ \Delta_i \in\ \|T \wedge \phi\| \\ T \qquad if\ \|T \wedge \phi\| = \emptyset \end{cases} \tag{14}$$

Example 5 *Consider again the background knowledge K_D defined in (1). Given that the contingential knowledge the agent knows is $E = \{cs(b)\}$ suppose that the new evidence $\phi = \{int(b, lp)\}$ is observed. The abductive framework P_A obtained in this situation is equal to:*

$$P_A = P_D \cup \left\langle \begin{array}{c} \{cs(b) \leftarrow\}, \left\{ \begin{array}{c} cs(b), int(b, ai), \\ int(b, pr\_cl) \end{array} \right\}; \\ \{\bot \leftarrow not\ int(b, lp)\} \end{array} \right\rangle$$

In this case the abductive answer is $\{\{int(b, ai)\}\}$. In agreement with our approach we would conclude that $T^{\diamond}_{int(b,lp)}$ is equal to:

$$\{T_1 = (K_D, cs(b) \cup int(b, ai) \cup int(b, lp))\} \tag{15}$$

We can see that the conjuction of P_D (the ζ-translation of K_D) with $\{cs(b) \leftarrow\ , int(b, ai) \leftarrow\ , int(b, lp) \leftarrow\ \}$ derives the following belief set:

$$BS = \{cs(b), int(b, lp), \neg int(b, lin), int(b, ai), ab0\} \tag{16}$$

which corresponds to advocated intuition.

Concluding Remarks

In this paper we present a new operator for the implicit belief revision of an agent. This operator is based on the concept that says that in many situations an agent should find an explanation for a recently observed fact before revising its beliefs. As a consequence, the proposed belief revision model does not observe the principle of minimal change. A translation of an agent's knowledge base, expressed by means of a conditional defaults set into an abductive extended logic program, is also presented. This translation allows us to make explicit the specificity that exists among defaults and also the logic program allows us to determine the possible explanations for a recently observed fact.

Through this abductive framework we can model the incorporating of new evidence into agent's belief base. It is done considering that the defaults exist as conditional declarations, which can be blocked through the abduction of more specific premisses that sustain new evidence, and defeat all counterarguments that exist in agent's known context. Therefore we say that the abductible answers are predictive explanations that nonmonotonically entails the new observation.

However a number of questions remain to be explored. In particular we are currently investigating how our model might be extended to incorporate plausibility degree to each possible explanation (Garcia & Lopes 1998). So that the most plausible explanation, in each context, is preferred. The incorporation of plausibility measures enables agents to weight the convenience of maintaining its beliefs instead of taking into consideration new evidence, thus enable an agent to behave pro-actively in order to clarify the situation.

All the examples that were presented in this paper were experimented with the latest version of the program REVISE (Damásio, Pereira, & Schroeder 1996); an extended logic programming system for revising knowledge bases. This program is based in top-down deerivation procedures for WFSX (Well Founded Semantics with eXplicit negation) (Alferes & Pereira 1996).

Acknowledgments

The first author was supported by CNPq - Brazil and Departamento de Informática, Universidade Federal do Espírito Santo (UFES) - Brazil. This work has been done in the framework of project DIXIT (2/2.1/TIT/1670/95) funded by programme Praxis XXI.

References

Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: partial meet functions for contraction and revision. *Journal of Symbolic Logic* 50:510–530.

Alferes, J. J., and Pereira, L. M. 1996. *Reasoning with Logic Programming.* Springer-Verlag: Heidelberg, Germany.

Boutilier, C. 1991. *Conditional Logics for Default Reasoning and Belief Revision.* Ph.D. Dissertation, Department of Computer Science, University of Toronto.

Boutilier, C. 1997. A unified model of qualitative belief change: A dynamical systems perspective. *Artificial Intelligence.* Forthcoming.

Damásio, C. V.; Pereira, L. M.; and Schroeder, M. 1996. Revise progress report. In *Workshop on Automated Reasoning: Bridging the Gap between Theory and Practice.*

Eshghi, K., and Kowalski, R. A. 1989. Abduction compared with negation by failure. In Levi, Giorgio; Martelli, M., ed., *Proceedings of the 6th International Conference on Logic Programming (ICLP '89)*, 234-254. Lisbon, Portugal: MIT Press.

Friedman, N. 1997. *Modeling Beliefs in Dynamic Systems.* Ph.D. Dissertation, Department of Computer Science, University of Stanford.

Garcia, B. B., and Lopes, G. P. 1998. Introducing plausibility measures to the process of belief revision through extended logic programs. Forthcoming.

Gärdenfors, P. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* Cambridge, MA: MIT Press, Bradford Books.

Ginsberg, M. L. 1986. Counterfactuals. *Artificial Intelligence* 30:35–79.

Goldszmidt, M., and Pearl, J. 1996. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence* 84(1-2):57–112.

Katsuno, H., and Mendelzon, A. O. 1991. On the difference between updating a knowledge base and revising it. In Allen, J.; Fikes, R.; and Sandewall, E., eds., *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, 387–394. San Mateo, CA, USA: Morgan Kaufmann Publishers.

Li, R., and Pereira, L. M. 1996. Knowledge assimilation in domains of actions: A possible causes approach. *Journal of Applied Non-Classical Logic, Special issue on Inconsistency Handling in Knowledge Systems.*

Pearl, J. 1990. System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning. *TARK: Theoretical Aspects of Reasoning about Knowledge* 3.

Winslett, M. 1988. Reasoning about action using a possible models approach. In Smith, Reid G.; Mitchell, T. M., ed., *Proceedings of the 7th National Conference on Artificial Intelligence*, 89–93. St. Paul, MN: Morgan Kaufmann.