Labeled Logics of Conditional Goals

Leendert W.N. van der Torre[1]

In the usual approaches to planning in AI, a planning agent is provided with a description of some state of affairs, a *goal state*, and charged with the task of discovering (or performing) some sequence of actions to achieve that goal. In particular, goals serve as computationally useful heuristic approximations of the relative preferences over the possible results of a plan. For example, goals are used to communicate desires in a compact and efficient way [5]. In the context of *qualitative* decision theory recently several logics for conditional or context-sensitive goals and desires have been proposed [3, 2, 8, 1, 10, 9, 6]. In this paper we introduce a version of a labeled deductive system as it was introduced by Gabbay in [4] to reason about goals. It has some desirable properties not found in other proposals.

First, the labeled logics formalize that goals interact and conflict. Goals only impose partial preferences, i.e. preferences given some objective and given some context. As a consequence, goals with overlapping contexts can conflict, because objectives can conflict. For example, to minimize time a tank has to be filled quickly, but to minimize loss it must be filled slowly. This cannot easily be formalized in standard formalisms. For example, the following counterintuitive derivation has to be blocked, where $G(\alpha)$ is read as 'preferably α.'

$$\dfrac{\dfrac{G(p)}{G(p \vee q)} \quad G(\neg p)}{\dfrac{G(q \wedge \neg p)}{G(q)}}$$

Second, the labeled logics are stronger than previous proposals in the sense that they validate strengthening of the antecedent (SA) and transitivity (TRANS). They can *only* be combined with the desirable inference pattern weakening of the consequent (WC) if additional machinery like labels is introduced in the logic. For example, the following counterintuitive derivation has to be blocked, where $G(\alpha|\beta)$ is read as 'preferably α if β' and $\top$ stands for a tautology.

$$\dfrac{\dfrac{G(c|\top)}{G(c \vee t|\top)}\ \text{WC}}{G(c \vee t|\neg c)}\ \text{SA}$$

Labeled goals $G(\alpha\,|\beta)_L$ can *roughly* be read as 'preferably α if β, against the background of L.' The label formalizes the context in which the goal is derived. It consists of the fulfillments (F) and violations (V) of the premises from which the goal is derived, where F and V are sets of propositional formulas. Given a premise $G(\alpha|\beta)$, its fulfillment is $\alpha \wedge \beta$ and its violation is $\neg\alpha \wedge \beta$ [12]. In the labeled logics of goals we use different types of consistency checks to restrict derivations. For example, the following derivation illustrates that a consistency check of the antecedent with F (the first set of the label) blocks the derivation above.

$$\dfrac{\dfrac{G(c|\top)_{(\{c\},\{\neg c\})}}{G(c \vee t|\top)_{(\{c\},\{\neg c\})}}\ \text{WC}}{G(c \vee t|\neg c)_{(\{c\},\{\neg c\})}}\ \text{SA}$$

After the introduction of the labeled logic we discuss transitivity and the disjunction rule. They formalize respectively that rules can be applied one after the other and reasoning by cases. Finally we discuss related research in deontic logic.

1 Definitions and examples

In this section we introduce a version of a labeled deductive system as it was introduced by Gabbay in [4]. Roughly speaking, the label L of a goal $G(\alpha|\beta)_L$ is a record of the fulfillments and violations of the premises that are used in the derivation of $G(\alpha|\beta)$.

Definition 1 (Language) *The language of* LLG *is a propositional base logic $\mathcal{L}$ and labeled dyadic goals $G(\alpha|\beta)_L$, with α and β sentences of $\mathcal{L}$, and L a pair (F, V) of sets of sentences of $\mathcal{L}$ (fulfillments and violations). We write $\models$ for entailment in $\mathcal{L}$.*

Each formula occurring as a premise has a label that consists of its own fulfillment and violation.

Definition 2 (Premise) *A formula $G(\alpha\,|\,\beta)_{(\{\alpha\wedge\beta\},\{\neg\alpha\wedge\beta\})}$ is called a premise of* LLG *when $\alpha \wedge \beta$ is consistent in $\mathcal{L}$.*

The label of a goal derived by an inference rule is the union of the labels of the premises used in this inference rule.

Definition 3 (LLG) *The logic* LLG *based on a violation check and a fulfillment check consists of inference rules (see top next page) extended with a condition $R = R_V + R_F$.*

R_V**:** *$G(\alpha\,|\beta)_{(F,V)}$ may only be derived if $\alpha \wedge \beta \not\models \gamma$ for all $\gamma \in V$: fulfilling a derived goal should not imply a violation of one of the goals it is derived from;*

R_F**:** *$G(\alpha\,|\beta)_{(F,V)}$ may only be derived if $\alpha \wedge \beta \not\models \neg\gamma$ for all $\gamma \in F$: it must always be possible to fulfill a derived goal and each of the goals it is derived from, though not necessarily all of them at the same time.*

[1] IRIT, Porte 324, University Paul Sabatier, 118 Route de Narbonne, 31062 Toulouse, France, torre@irit.fr

© 1998 L.W.N. van der Torre
ECAI 98. *13th European Conference on Artificial Intelligence*
Young Researcher Paper
Edited by Henri Prade
Published in 1998 by John Wiley & Sons, Ltd.

The inference rules of LLG *are replacements by logical equivalents and the following five rules.*

$$\text{SA}_R : \frac{G(\alpha \mid \beta_1)_{(F,V)}, R}{G(\alpha \mid \beta_1 \wedge \beta_2)_{(F,V)}}$$

$$\text{WC}_R : \frac{G(\alpha_1 \mid \beta)_{(F,V)}, R}{G(\alpha_1 \vee \alpha_2 \mid \beta)_{(F,V)}}$$

$$\text{TRANS}_R : \frac{G(\alpha|\beta \wedge \gamma)_{(F_1,V_1)}, G(\beta|\gamma)_{(F_2,V_2)}, R}{G(\alpha \wedge \beta \mid \gamma)_{(F_1 \cup F_2, V_1 \cup V_2)}}$$

$$\text{AND}_R : \frac{G(\alpha_1 \mid \beta)_{(F_1,V_1)}, G(\alpha_2 \mid \beta)_{(F_2,V_2)}, R}{G(\alpha_1 \wedge \alpha_2 \mid \beta)_{(F_1 \cup F_2, V_1 \cup V_2)}}$$

$$\text{OR}_R : \frac{G(\alpha \mid \beta_1)_{(F_1,V_1)}, G(\alpha \mid \beta_2)_{(F_2,V_2)}, R}{G(\alpha \mid \beta_1 \vee \beta_2)_{(F_1 \cup F_2, V_1 \cup V_2)}}$$

We say $\{G(\alpha_i \,|\, \beta_i) \mid 1 \leq i \leq n\} \vdash_{\text{LLG}} G(\alpha \,|\, \beta)$ *if there is a labeled goal* $G(\alpha|\beta)_L$ *that can be derived from the set of goals* $\{G(\alpha_i|\beta_i)_{(\{\alpha_i \wedge \beta_i\}, \{\neg\alpha_i \wedge \beta_i\})} \mid 1 \leq i \leq n\}$.

It is easily checked that conflicting goals are consistent, and that the counterintuitive derivations in the introduction are blocked. The following example illustrates that the transitivity rule formalizes that one conditional rule after the other can be applied. Derivations go 'as far as possible.'

Example 1 *Consider the following set of goals.*

$$S = \{G(a|b), G(b|c), G(c|\top), G(\neg a|\top)\}$$

There is a conflict for a*, because we have* $S \vdash_{\text{LLG}} G(a|\top)$ *and* $S \vdash_{\text{LLG}} G(\neg a|\top)$*. There is not a conflict for* b*, because we have* $S \vdash_{\text{LLG}} G(b|\top)$ *and* $S \not\vdash_{\text{LLG}} G(\neg b|\top)$*. Hence, derivation chains go as far as possible. Moreover, the following derivation illustrates how more complex forms of transitivity are supported.*

$$\dfrac{\dfrac{G(a|b)_{(\{a\wedge b\},\{\neg a\wedge b\})}}{G(a|b\wedge\neg c)_{(\{a\wedge b\}),\{\neg a\wedge b\})}}\,\text{SA} \quad G(\neg c|b\vee c)_{(\{b\wedge\neg c\},\{c\})}}{\dfrac{G(a\wedge b\wedge\neg c|b\vee c)_{(\{a\wedge b, b\wedge\neg c\},\{\neg a\wedge b, c\})}}{G(a|b\vee c)_{(\{a\wedge b, b\wedge\neg c\},\{\neg a\wedge b, c\})}}\,\text{WC}}\,\text{TRANS}$$

Our second example illustrates how the disjunction rule supports reasoning by cases, which is usually considered intuitive in e.g. conditional logic and default logic.[2]

Example 2 *We have* $G(a\,|\,b), G(a\,|\,\neg b) \not\vdash_{\text{LLG}} G(a\,|\,b \leftrightarrow a)$ *as a result of the fulfillment check, as can be verified with the following derivation.*

$$\dfrac{\dfrac{G(a|b)_{(\{a\wedge b\},\{\neg a\wedge b\})} \quad G(a|\neg b)_{(\{a\wedge\neg b\},\{\neg a\wedge\neg b\})}}{G(a|\top)_{(\{a\wedge b, a\wedge\neg b\},\{\neg a\wedge b, \neg a\wedge\neg b\})}}\,\text{OR}}{G(a|b\leftrightarrow a)_{(\{a\wedge b, a\wedge\neg b\},\{\neg a\wedge b, \neg a\wedge\neg b\})}}\,\text{SA}$$

The following derivation illustrates a more complex type of reasoning by cases.

$$\dfrac{\dfrac{G(a\wedge c|b)_{(\{a\wedge c\wedge b\},\{\neg(a\wedge c)\wedge b\})}}{G(a|b)_{(\{a\wedge c\wedge b\},\{\neg(a\wedge c)\wedge b\})}} \quad \dfrac{G(a\wedge\neg c|\neg b)_{(\{a\wedge\neg c\wedge\neg b\},\{\neg(a\wedge\neg c)\wedge\neg b\})}}{G(a|\neg b)_{(\{a\wedge\neg c\wedge\neg b\},\{\neg(a\wedge\neg c)\wedge\neg b\})}}}{G(a|\top)_{(\{a\wedge c\wedge b, a\wedge\neg c\wedge\neg b\}),\{\neg(a\wedge c)\wedge b, \neg(a\wedge\neg c)\wedge\neg b\})}}\,\text{OR}$$

[2] However, it can be questioned whether the disjunction rule of the antecedent formalizes reasoning by cases satisfactorily. In particular, it might be argued that there is an implicit modal operator that falsifies the disjunction rule, e.g. when $G(\alpha\,|\,\beta)$ stands for 'preferably α if it is unalterable / known / believed that β.'

2 Closing remarks

Several authors have observed the relation between goals and desires in qualitative decision theory and obligations in deontic logic [8, 1, 6] and we have discussed the relation between decision theory, diagnosis theory and deontic logic in [12]. We argued that decision theory not only represent violations, like diagnosis theory, but also fulfillments, because it not only reasons about the past (with incomplete knowledge), but also about the future (with uncertain knowledge). The logics introduced in this paper are refinements of the labeled deontic logic proposed in [13, 14, 11] with these distinctions, such that the disjunction rule can be supported (and therefore also reasoning by cases). Makinson [7] introduced an alternative labeled deontic logic that is closely related to our logics.

The logic can be extended with the following additional completeness assumption: *a goal is completely defined by its fulfillment and violation condition.* This means that goals can be defined by $G(\alpha\,|\,\beta) =_{def} \alpha \wedge \beta \succ \neg\alpha \wedge \beta$ where $\succ$ can be read as an operator from preference logic. Hence, $G(\alpha\,|\,\beta)$ is true if $\alpha \wedge \beta$ is preferred to $\neg\alpha \wedge \beta$. However, this definition implies the theorem $G(\alpha|\beta \wedge \gamma) \leftrightarrow G(\alpha \wedge \beta|\beta \wedge \gamma)$ which is counterintuitive at first sight.

REFERENCES

[1] C. Boutilier, 'Toward a logic for qualitative decision theory', in *Proceedings of the KR'94*, pp. 75–86, (1994).

[2] J. Doyle, Y. Shoham, and M.P. Wellman, 'The logic of relative desires', in *Sixth International Symposium on Methodologies for Intelligent Systems*, Charlotte, North Carolina, (1991).

[3] J. Doyle and M.P. Wellman, 'Preferential semantics for goals', in *Proceedings of AAAI-91*, pp. 698–703, Anaheim, (1991).

[4] D. Gabbay, *Labelled Deductive Systems*, volume 1, Oxford University Press, 1996.

[5] P. Haddawy and S. Hanks, 'Representations for decision-theoretic planning: Utility functions for dead-line goals', in *Proceedings of the KR'92*, Cambridge, MA, (1992).

[6] J. Lang, 'Conditional desires and utilities - an alternative approach to qualitative decision theory', in *Proceedings of the ECAI'96*, pp. 318–322, (1996).

[7] D. Makinson, 'The fundamental problem of deontic reasoning', in *Proceedings of the Fourth International Workshop on Deontic Logic in Computer Science (ΔEON'98)*, pp. 3–42, (1998).

[8] J. Pearl, 'From conditional ought to qualitative decision theory', in *Proceedings of the UAI'93*, pp. 12–20, (1993).

[9] S.-W. Tan and J. Pearl, 'Qualitative decision theory', in *Proceedings of the AAAI'94*, (1994).

[10] S.-W. Tan and J. Pearl, 'Specification and evaluation of preferences under uncertainty', in *Proceedings of the KR'94*, pp. 530–539, (1994).

[11] L.W.N. van der Torre, *Reasoning About Obligations: Defeasibility in Preference-based Deontic Logic*, Ph.D. dissertation, Erasmus University Rotterdam, 1997.

[12] L.W.N. van der Torre and Y.-H. Tan, 'Diagnosis and decision making in normative reasoning', *Artificial Intelligence and Law*, (1998). To appear.

[13] L.W.N. van der Torre and Y.H. Tan, 'Cancelling and overshadowing: two types of defeasibility in defeasible deontic logic', in *Proceedings of the IJCAI'95*, pp. 1525–1532. Morgan Kaufman, (1995).

[14] L.W.N. van der Torre and Y.H. Tan, 'The many faces of defeasibility in defeasible deontic logic', in *Defeasible Deontic Logic*, ed., D. Nute, 79–121, Kluwer, (1997).