ELSEVIER Artificial Intelligence 94 (1997) 269–270

Artificial Intelligence

Forthcoming Papers

F. Kabanza, M. Barbeau and R. St-Denis, Planning control rules for reactive agents

A traditional approach for planning is to evaluate goal statements over state trajectories modeling predicted behaviors of an agent. This paper describes a powerful extension of this approach for handling complex goals for reactive agents. We describe goals by using a modal temporal logic that can express quite complex time, safety, and liveness constraints. Our method is based on an incremental planner algorithm that generates a reactive plan by computing a sequence of partially satisfactory reactive plans converging to a completely satisfactory one. Partial satisfaction means that an agent controlled by the plan accomplishes its goal only for some environment events. Complete satisfaction means that the agent accomplishes its goal whatever environment events occur during the execution of the plan. As such, our planner can be stopped at any time to yield a useful plan. An implemented prototype is used to evaluate our planner on empirical problems.

T.C. Przymusinski, Autoepistemic logic of knowledge and beliefs

In recent years, various formalizations of non-monotonic reasoning and different semantics for normal and disjunctive logic programs have been proposed, including autoepistemic logic, circumscription, *CWA*, *GCWA*, *ECWA*, epistemic specifications, stable, well-founded, stationary and static semantics of normal and disjunctive logic programs.

In this paper we introduce a simple non-monotonic knowledge representation framework which isomorphically contains all of the above mentioned non-monotonic formalisms and semantics as special cases and yet is significantly more expressive than each one of these formalisms considered individually. The new formalism, called the *Autoepistemic Logic of Knowledge and Beliefs*, *AELB*, is obtained by augmenting Moore's autoepistemic logic, *AEL*, already employing the *knowledge operator*, $\mathcal{L}$, with an additional *belief operator*, $\mathcal{B}$. As a result, we are able to reason not only about formulae F which are *known* to be true (i.e., those for which $\mathcal{L}F$ holds) but also about those which are only *believed* to be true (i.e., those for which $\mathcal{B}F$ holds).

The proposed logic constitutes a powerful new formalism which can serve as a *unifying framework* for several major non-monotonic formalisms. It allows us to better understand mutual relationships existing between different formalisms and semantics and enables us to provide them with simpler and more natural definitions. It also naturally leads to new, even more expressive, flexible and modular formalizations and semantics.

T.L. McCluskey and J.M. Porteous, Engineering and compiling planning domain models to promote validity and efficiency

This paper postulates a rigorous method for the construction of classical planning domain models. We describe, with the help of a non-trivial example, a tool supported method for encoding such models. The method results

Elsevier Science B.V.

in an "object-centred" specification of the domain that lifts the representation from the level of the literal to the level of the object. Thus, for example, operators are defined in terms of how they change the state of objects, and planning states are defined as amalgams of the objects' states. The method features two classes of tools: for initial capture and validation of the domain model; and for operationalising the domain model (a process we call compilation) for later planning. Here we focus on compilation tools used to generate macros and goal orders to be utilised at plan generation time. We describe them in depth, and evaluate empirically their combined benefits in plan-generation speed-up.

The method's main benefit is in helping the modeller to produce a tight, valid and operational domain model. It also has the potential benefits of (i) forcing a change of emphasis in classical planning research to encompass knowledge-based aspects of target planning domains in a systematic manner, (ii) helping to bridge the gap between the research area of theoretical but unrealistic planning on the one hand, and "scruffy" but real-world planning on the other, (iii) a commitment to a knowledge representation form which allows powerful techniques for planning domain model validation and planning algorithm speed-up can be bound up into a tool-supported environment.

A. Darwiche, A logical notion of conditional independence: properties and applications

D. Galles and J. Pearl, Axioms of causal relevance

R. Greiner, A.J. Grove and A. Kogan, Knowing what doesn't matter: exploiting the omission of irrelevant data

R. Khardon and D. Roth, Defaults and relevance in model based reasoning

J. Kivinen, M.K. Warmuth and P. Auer, The Perceptron algorithm versus Winnow: linear versus logarithmic mistake bounds when few input variables are relevant

R. Kohavi and G.H. John, Wrappers for feature subset selection

G. Lakemeyer, Relevance from an epistemic perspective

A.Y. Levy, R.E. Fikes and Y. Sagiv, Speeding up inferences using relevance reasoning: a formalism and algorithms

C.A. Gunter, Teow-Hin Ngair and D. Subramanian, The common order-theoretic structure of version spaces and ATMSs

Y. Xia, S.S. Iyengar and N.E. Brener, An event driven integration reasoning scheme for handling dynamic threats in an unstructured environment (Research Note)

N. Friedman and J.Y. Halpern, Modeling belief in dynamic systems, Part 1: Dynamics

E. Giunchiglia, G. Neelakantan Kartha and V. Lifschitz, Representing action: indeterminacy and ramifications