DOWNLOAD PDF AUTOEPISTEMIC LOGIC

Chapter 1 : Epistemic Logic (Stanford Encyclopedia of Philosophy)

The autoepistemic logic is a formal logic for the representation and reasoning of knowledge about knowledge. While propositional logic can only express facts, autoepistemic logic can express knowledge and lack of knowledge about facts.

Show Context Citation Context For a nonground program P , the stable models are those of the ground instantiation of the program with respect to its Herbrand universe. The constraint interpretation of the rules provided by the s Stable models and an alternative logic programming paradigm by Victor W. In this paper we reexamine the place and role of stable model semantics in logic programming and contrast it with a least Herbrand model approach to Horn programs. We demonstrate that inherent features of stable model semantics naturally lead to a logic programming system that offers an interesting We demonstrate that inherent features of stable model semantics naturally lead to a logic programming system that offers an interesting alternative to more traditional logic programming styles of Horn logic programming, stratified logic programming and logic programming with well-founded semantics. The proposed approach is based on the interpretation of program clauses as constraints. In this setting programs do not describe a single intended model, but a family of stable models. These stable models encode solutions to the constraint satisfaction problem described by the program. Our approach imposes restrictions on the syntax of logic programs. In particular, function symbols are eliminated from the language. We argue that the resulting logic programming system is well-attuned to problems in the class NP, has a well-defined domain of applications, and an emerging methodology of programming. We point out that what makes the whole approach viable is recent progress in implementations of algorithms to compute stable models of propositional logic programs. Warren , " SLD resolution with negation as finite failure SLDNF reflects the procedural interpretation of predicate calculus as a programming language and forms the computational basis for Prolog systems. Despite its advantages for stack-based memory management, SLDNF is often not appropriate for query evalu Despite its advantages for stack-based memory management, SLDNF is often not appropriate for query evaluation for three reasons: We address three problems fir a goal-oriented query evaluation of general logic programs by presenting tabled evaluation with delaying SLG resolution. The alternating fixpoint of a logic program with negation is defined constructively. The underlying idea is monotonically to build up a set of negative conclusions until the least fixpoint is reached, using a transformation related to the one that defines stable models. From a fixed set of negative From a fixed set of negative conclusions, the positive conclusions follow without deriving any further negative ones , by traditional Horn clause semantics. The union of positive and negative conclusions is called the alternating xpoint partial model. The principal contributions of this work are 1 that the alternating fixpoint partial model is identical to the well-founded partial model, and 2 that alternating xpoint logic is at least as expressive as xpoint logic on all structures. Also, on finite structures, fixpoint logic is as expressive as alternating fixpoint logic. In this paper, we review recent work aimed at the application of declarative logic programming to knowledge representation in artificial intelligence. We consider exten- sions of the language of definite logic programs by classical strong negation, disjunc- tion, and some modal operators and sh We consider exten- sions of the language of definite logic programs by classical strong negation, disjunc- tion, and some modal operators and show how each of the added features extends the representational power of the language. Przymusinski - New Generation Computing , " We introduce the stable model semantics for disjunctive logic programs and deductive databases, which generalizes the stable model semantics, defined earlier for normal i. Depending on whether only total 2-valued or all partial 3-valued models are used we obtain th Depending on whether only total 2-valued or all partial 3-valued models are used we obtain the disjunctive stable semantics or the partial disjunctive stable semantics, respectively. The proposed semantics are shown to have the following properties:

Chapter 2 : Talk:Autoepistemic logic - Wikipedia

Autoepistemic logic is one of the principal modes of nonmonotonic reasoning. It unifies several other modes of nonmonotonic reasoning and has important application in logic programming. In the paper, a theory of autoepistemic logic is developed.

Cautious Monotony is the converse of Cut: Taken together these principles express that inference is a cumulative enterprise: Accordingly, these properties are commonly taken to be central principles of NML. Rational Monotony is a more controversial property than Cautious Monotony. For instance, Stalnaker gave a counter-example to Rational Monotony see the supplementary document in view of which it is clear that in not all application contexts Rational Monotony is a desirable property of defeasible reasoning. Dealing with conflicts A separate issue from the formal properties of a non-monotonic consequence relation, although one that is strictly intertwined with it, is the issue of how conflicts between potential defeasible conclusions are to be handled. We can distinguish two types of conflict handling in defeasible reasoning both of which will be discussed in more detail below. On the one hand, we have conflict resolution principles such as the Specificity Principle or other measures of argument strength. On the other hand, we have reasoning types that deal in different ways with non-resolvable conflicts, most prominently the Credulous and the Skeptical reasoning types. In this section we still stay on an abstract level while concrete NMLs are discussed in the section Non-Monotonic Formalisms. When a conflict of either kind arises, steps have to be taken to preserve or restore consistency. Conflicts of type i are to be resolved in favor of the hard facts in the sense that the conflicting defeasible conclusion is to be retracted. More interesting are mechanisms to resolve conflicts of type ii. In order to analyze these, we will make use of schematic inference graphs similar to the ones used in Inheritance Networks , see below. For instance, our previous example featuring the bird Tweety is illustrated as follows: We use the following conventions: So, we can read the diagram as follows: We have a conflict between the following two arguments where arguments are sequences of inferences: Both arguments include a final defeasible inference. According to the Specificity Principle an inference with a more specific antecedent overrides a conflicting defeasible inference with a less specific antecedent. Logicians distinguish between strong and weak specificity: Note that the difference concerns the nature of the link between A and B. In the context of legal reasoning we may have the principles lex superior resp. We give some examples. For a systematic survey and classification of preference handling mechanisms in NML the interested reader is referred to Delgrande et al. According to the Weakest Link Principle Pollock an argument is preferred over another conflicting argument if its weakest defeasible link is stronger than the weakest defeasible link in the conflicting argument. Take, for example, the situation in the following figure: On the left we see an inference graph with two conflicting arguments. On the right we see the preference ordering. Another approach to preferences is procedural. Roughly, it instructs to always apply the rule with the highest priority first. There may be various such rules with highest priority, but for the sake of simplicity we neglect this possibility in what follows. Take the following example that is frequently discussed in the literature. So we derive B. Brewka and Eiter argue against the procedural approach and that it is more intuitive to derive B and C. Delgrande and Schaub argue that the example presents an incoherent set of rules. This is put in question in Horty where a consistent deontic reading in terms of conditional imperatives is presented which also challenges the procedural approach by favoring the conclusions B and C. Ford pointed out that the order of strict and defeasible links in arguments matters. The reason is that in the former case it is not possible that no A is a D while in the second case it is possible that no A is a not-D. This is illustrated in the following figure: Whenever the extension of A is non-empty there will be As that are Ds. These two options correspond to significantly different ways to construe a given body of defeasible knowledge, and yield different results as to what defeasible conclusions are warranted on the basis of such a knowledge base. The difference between these basic attitudes comes to this. In the presence of potentially conflicting defeasible inferences and in the absence of further considerations such as specificity â€” see above , the credulous reasoner always commits to as many defeasible conclusions as possible, subject to a consistency requirement, whereas the skeptical reasoner withholds assent from potentially conflicted defeasible conclusions. Suppose our knowledge base contains defeasible information to the effect that a given individual, Nixon, is both a Quaker and a Republican. Quakers, by and large, are pacifists, whereas Republicans, by and large, are not. The question is what defeasible conclusions are warranted on the basis of this body of knowledge, and in particular whether we should infer that Nixon is a pacifist or that he is not a pacifist. While the credulous reasoner commits to both conclusions, the skeptical reasoner refrains from either. A rationale behind the credulous reasoning type is to provide an overview of possible conclusions given the conflicting defeasible inferences in order to subsequently make a choice among them. This is especially interesting in practical reasoning contexts in which the choice determines a course of action and in which extra-logical considerations based on preferences, values, etc. In contrast, the rationale behind the skeptical reasoning type is to determine uncontested defeasible conclusions. The first issue is illustrated in the following figure: Since penguins do not fly, we know that penguins are exceptional birds, at least with respect to the property of flying. A very cautious reasoner may take this to be a reason to be skeptical about attributing other typical properties of birds to penguins. NMLs that do not allow to derive has wings are said to suffer from the Drowning Problem or to not satisfy the Strong Independence property. The question whether the exceptional status of Penguin relative to flies should spread also to other properties of Bird may depend on specific relevance relations among these properties. For instance, Koons proposes that causal relations play a role: Similarly, Pelletier and Elio argue that explanatory relations play a significant role in the way in which reasoners treat exceptional information in nonmonotonic inference. Another much discussed issue e. Such conclusions are called Floating Conclusions. The following figure illustrates this with an extended version of the Nixon Diamond. We conclude our discussion with so-called Zombie-Arguments Makinson and Schlechta , Touretzky et al. Recall that a skeptical reasoner does not commit to a conflicting argument. Makinson and Schlechta argue that super-arguments of such conflicted arguments â€”although not acceptableâ€” nevertheless still have the power to undermine the commitment of a reasoner to an otherwise unconflicted argument. We see an example in the following figure: Non-monotonic formalisms Pioneering work in the field of NMLs began with the realization that in order to give a mathematically precise characterization of defeasible reasoning CL is inadequate. Such a realization was accompanied by the effort to reproduce the success of CL in the representation of mathematical, or formal, reasoning. Reiter see Ginsberg for a collection of early papers in the field and Gabbay et al. In , the Artificial Intelligence Journal published an issue vol. In this section an overview will be provided on some important formalisms. Since the evolutionary tree of NMLs has grown extraordinarily rich, we will restrict the focus on presenting the basic ideas behind some of the most influential and well-known approaches. Database theory was one of the earliest sources of such examples, especially as regards the closed world assumption. The agents queries the database and, not surprisingly, responds that there are no direct flights. How does the travel agent know? What is at work here is a tacit assumption that the database is complete, and that since the database does not list any direct flights between the two cities, there are none. A useful way to look at this process is as a kind of minimization, i. Moreover, on pain of inconsistencies, such a minimization needs to take place not with respect to what the database explicitly contains but with respect to what it implies. Circumscription makes explicit the intuition that, all other things being equal, extensions of certain predicates should be minimal. Implicit in this principle is the idea that specimens should not be considered to be abnormal unless there is positive information to that effect. This basic form of circumscription has been generalized, for, in practice, one needs to minimize the extension of a predicate, while allowing the extension of certain other predicates to vary. From the point of view of applications, however, circumscription has a major computational shortcoming, which is due to the nature of the second-order language in which circumscription is formulated. Another influential mechanism realizing the closed world assumption is Negation as Failure or Default Negation. It can nicely be explained if we take a look at Logic Programming. A logic program consists of a list of rules such as: This has been generalized in various ways e. A concrete example for a rule is: Default negation is realized as follows: The latter two rules will not be triggered, and the first rule be will interpreted with the closed world assumption so that flies can be inferred. Now suppose we know penguin. In this case the first rule is not applicable but the latter two are and we derive bird and not-flies.

However, there can be exceptions, which can interact in complex ways as in the following example: Things can be more complicated still, for in turn, baby bats are exceptional bats, in that they do not fly does that make them unexceptional mammals? Here we have potentially conflicting inferences. When we infer that Stellaluna, being a baby bat, does not fly, we are resolving all these potential conflicts based on the Specificity Principle. Non-monotonic inheritance networks were developed for the purpose of capturing taxonomic examples such as the one above. When exceptions are allowed, the network is interpreted defeasibly. The following figure gives a network representing this state of affair: More specific information overrides more generic information. Research on non-monotonic inheritance focuses on the different ways in which one can make this idea precise.

Chapter 3 : Autoepistemic logic - Wikipedia

Autoepistemic logic's wiki: The autoepistemic logic is a formal logic for the representation and reasoning of knowledge about knowledge. While propositional logic can only express facts, autoepistemic logic can express knowledge and lack of knowledge about facts.

The logics range from S4 over the intermediate systems S4. By way of example, Hintikka settled for S4 , Kutschera argued for S4. Similar completeness cataloguing is available for belief where axiom T is dropped and usually replaced by D to avoid the condition of truth for belief but retain consistency among beliefs yieding systems like KD4â€"KD45 for belief. This also paves the way for combining epistemic and doxastic systems and for studying the interplay between knowledge and belief see Voorbraak Care should be taken however not to collapse knowledge and belief in the combined systems as have been noted by Lenzen and Stalnaker , among others. A particularly malignant philosophical problem for epistemic logic is related to closure properties. There are various ways of dealing with logical omniscience. Some of the first proposals for solving the problem of logical omniscience introduce semantical entities which explain why the agent appears to be, but in fact is not really guilty of logical omniscience. The basic idea is that an agent may mistakenly count among the worlds consistent with his or her knowledge, some worlds containing logical contradictions. The mistake is simply a product of limited resources; the agent may not be in a position to detect the contradiction and may erroneously count them as genuine possibilities. Allowing impossible possible worlds or seemingly possible worlds in which the semantic valuation of the formulas is arbitrary to a certain extent provides a way of making the appearance of logical omniscience less threatening. After all, on any realistic account of epistemic agency, the agent is likely to consider albeit inadvertantly worlds in which the laws of logic do not hold. Since no real epistemic principles hold broadly enough to encompass impossible and seemingly possible worlds, some conditions must be applied to epistemic models such that they cohere with epistemic principles. Computer scientists have proposed that what is being modelled in epistemic logic is not knowledge simpliciter but a related concept which is immune to logical omniscience. The agents neither have to compute knowledge nor can they be held responsible for answering queries based on their knowledge under the implicit understanding of knowledge. Logical omniscience is an epistemological condition for implicit knowledge, but the agent may actually fail to realize this condition. Groups of Knowers Single-agent systems may be extended to groups or multi-agent systems. The primary difference between the semantics given for a mono-agent and a multi-agent semantics is roughly that n accessibility relations are introduced. A modal system for n agents is obtained by joining together n modal logics where for simplicity it may be assumed that the agents are homogenous in the sense that they may all be described by the same logical system. An epistemic logic for n agents consists of n copies of a certain modal logic. In such an extended epistemic logic it is possible to express that some agent in the group knows a certain fact, that an agent knows that another agent knows a fact etc. It is possible to develop the logic even further: Not only may an agent know that another agent knows a fact, but they may all know this fact simultaneously. From here it is possible to express that everyone knows that everyone knows that everyone knows, thatâ€¦. That it is common knowledge. As Lewis noted in his book Convention a convention requires common knowledge among the agents that observe it. A variety of norms, social and linguistic practices, agent interactions and games presuppose common knowledge Aumann Once multiple agents have been added to the syntax, the language is augmented with an additional operator c. Well-formed formulas follow the standard recursive recipe with a few, but obvious, modifications taking into account the multiple agents. To semantically interpret n knowledge operators, binary accessibility relations Rn are defined over the set of possible worlds W. The relation must be flexible enough to express the relationship between individual and common knowledge. The idea is to let the accessibility relation for c be the transitive closure of the union of the accessibility relations corresponding to the singular knowledge operators. The semantics for the Boolean connectives remain intact. The formula KiA is true in a world w, i. The formula CA is true in a world w, i. Varying the properties of the accessibility relations R1,R2,â€¦,Rn, as described above results in different epistemic logics. For instance system K with common knowledge is determined by all frames, while system S4 with common knowledge is determined by all reflexive and transitive frames. Similar results can be obtained for the remaining epistemic logics Fagin et al. Active Agenthood A significant difference between alethic and epistemic logic is the introduction of the agent c to the syntax. But what role does the agent play in epistemic logic? At the early stages in the development of the logic they primarily served as indices on the accessibility relation between possible worlds. However, there is nothing particularly epistemic about being an index, and epistemic logicians soon began recognizing the central role of the agent much more explicitly. An agent may have knowledge which is S4. An important set of questions seem to be how the agent has to behave in order to gain the epistemic strength that he has. To make epistemic logic pertinent to epistemology, computer science, artificial intelligence and cognitive psychology the activity of agents must be included in our formal considerations. The original symbolic notation of a knowing agent also suggests this: An agent term should be inside the scope of the knowledge operator â€" not outside as Hintikka notes Inquiring agents are agents who read data, change their minds, interact or have common knowledge, act according to strategies and play games, have memory and act upon it, follow various methodological rules, expand, contract or revise their knowledge bases, etc. Inquiring agents are active agents Hendricks Game theory is about strategies for winning games in the context of other agents. Game theory has therefore played a prominent role in reflections on epistemic agency the study of the behavior of interactive epistemic agents. Aumann, van Benthem, Brandenburger, Fagin, Halpern, Keisler, Moses, Stalnaker, Vardi and others have demonstrated how logical epistemology uncovers important features of agent rationality showing how game theory adds to the general understanding of notions like knowledge, belief and belief revision. Mixing the theory of belief change and epistemic logic furnishes an illustrative example of active agents. The idea dates back to the mid s. In , de Rijke showed that the AGM-axioms governing expansion and revision may be translated into the object language of dynamic modal logic de Rijke At about the same time, Segerberg demonstrated how the entire theory of belief revision could be formulated in modal logic. Similarly for other combinations of the belief operator with negation. Active agenthood is also realizable directly on the agent level. One may also choose to endow the agents with epistemic capacities facilitating special epistemic behaviors. Perfect recall is in turn an epistemic recommendation telling the agent to remember his earlier epistemic states. There are other structural properties of agents being studied in the literature of dynamic epistemic logics. Multi-Modalities When epistemic logic was still in its infancy, Dana Scott noted: Here is what I consider one of the biggest mistakes of all in modal logic: The only way to have any philosophically significant results in deontic logic or epistemic logic is to combine these operators with: Tense operators otherwise how can you formulate principles of change? Scott , Nowadays, there are various ways in which multi-modalities may be realized in epistemic logic. One standard way is to follow Fagin et al. The run may be construed as an account of the behavior of the system for possible executions. For every time, the system is in some global state as a function of the particular time. The system may be thought of as a series of runs rather than agents. In this case, what is being modeled are the possible behaviors of the system over a collection of executions. A system of this kind defines a Kripke-structure with an equivalence relation over points. Truth of a formula is given with respect to a point. If truth is relative to a point then there is a question of when which opens up for the introduction of temporal operators. One may for instance define a universal future-tense operator such that a formula is true relative to the current point and all later points. The mixture of epistemic and temporal operators can handle claims about the temporal development of knowledge in the system. Epistemic Logic and Epistemology Generally speaking, contemporary epistemology is organized around two major goals: The first of these goals has, for the most part, been a concern of philosophers who rely on thought experiments, traditional conceptual analysis or intuitions-based methods of various kinds. By contrast, philosophers working with epistemic logic have pursued the second goal. The apparent divergence of both enterprises can be reconciled to some extent once one recognizes that both goals relate to a third, and possibly more general problem, namely 3 the problem of understanding the rationality of inquiry. This problem is of equal importance to both epistemic logicians and traditional epistemologists. Dynamical treatments of epistemic logic and insights into the logic of inquiry from epistemic logicians speak directly to this third, unifying goal. In recent decades, it is precisely the dynamical model of knowledge and inquiry that has concerned philosophically-inclined epistemic logicians. For a systematic treatment of the interplay between epistemology and epistemic logic we refer to Hendricks and Symons Monographs and Anthologies Boh, I. Epistemic Logic in the Middle Ages. Concurrent Dynamic Epistemic Logic. Handbook of Philosophical Logic, vol. Extensions of Classical Logic. A Textbook on Belief Revision. Kluwer Academic Publishers Hendricks, V. The Convergence of Scientific Knowledge--a view from the limit. Studia Logica Library Series, Dordrecht: F, and Pedersen, S. Mainstream and Formal Epistemology. New Waves in Epistemology. Ch and Hoek, W. Cambridge Tracts in Theoretical Computer Science An Introduction to the Logic of the Two Notions.

DOWNLOAD PDF AUTOEPISTEMIC LOGIC

Chapter 4 : CiteSeerX â€" Citation Query Autoepistemic logic programming

Autoepistemic logic is one of the principal modes of nonmonotonic reasoning. It unifies several other modes of nonmonotonic reasoning and has important application in logic programming. In the.

Query languages for retrieving information from disjunctive databases are an interesting open area of research. In this paper we study the expressive power of major non-monotonic formalisms -- such as circumscription, default logic, autoepistemic logic and some logic programming languages -- used as In this paper we study the expressive power of major non-monotonic formalisms -- such as circumscription, default logic, autoepistemic logic and some logic programming languages -- used as query languages over disjunctive databases. For this aim, we define the semantics of query expressions formulated in different nonmonotonic logics. The expressive power of the languages that we consider has been explored in the context of relational databases. Here, we extend this study to disjunctive databases; as a result, we obtain a finer-grained characterization of the expressiveness of those languages and interesting fragments thereof. For instance, we show that there exist simple queries that cannot be expressed by any preferential semantics including the minimal model semantics and the various forms of circumscription , while they can be expressed in default and autoepistemic logic Show Context Citation Context The main features of this calculus are simplicity and regularity, and the fact that proofs can be surprisingly concise and, in many cases, involve only a s The main features of this calculus are simplicity and regularity, and the fact that proofs can be surprisingly concise and, in many cases, involve only a small part of the default theory. On the other hand, the proof-theoretic aspects are not yet completely understood. The fundamental papers by Gabbay [14] , Makinson [24] and Kraus, Lehmann and Magidor [19] , focus their attention on general properties of non-monotonic inference, rather than on specific formalisms. In particular, they do not axio A Sequent Calculus for Circumscription by P. Olivetti , " This work is part of a larger project, aiming at a uniform proof-theoretic reconstruction of the major families of non-monotonic logics. Among the novelties of the calculus, we mention that CIRC is analytic Among the novelties of the calculus, we mention that CIRC is analytic and comprises an axiomatic rejection method, which allows for a fully detailed formalization of the nonmonotonic aspects of inference. The semantic and algorithmic aspects of non-monotonic reasoning have been extensively investigated e. The fundamental papers by Gabbay [13] , Makinson [23] and Kraus, Lehmann and Magidor [18] , focus their attention on general properties of The fundamental papers by Gabbay [13] , Makinson [23] and Kraus, Lehmann and Magidor [18] , focus their attention o Another section discusses and formalizes a sample heuristic rule for conjecturing whether an individual id Another section discusses and formalizes a sample heuristic rule for conjecturing whether an individual identity other than personal, being conveyed by a toponym, was used literally or fictitiously in a given historical corpus of legal casenotes. For example, a landlocked city being named and referred to as though it was a sea port is a fairly good cue for assuming that the toponym is a disguise. Yet, the interpretation is governed by other conventions, when in a play by Shakeaspeare it is stated that a given scene is set on the sea coast of Bohemia. Further discussion of a situational casuistry for identification especially individual and personal along with more formal representations will appear in a companion paper [84], also at the disciplinary meet of AI formalisms and legal applications. In Haas [34], a robot is made to reason ab

Chapter 5 : CiteSeerX â€" Citation Query Autoepistemic logic

Marjon Blondeel, Tommaso Flaminio, Steven Schockaert, LluÃ-s Godo, Martine De Cock, On the relationship between fuzzy autoepistemic logic and fuzzy modal logics of belief, Fuzzy Sets and Systems, v n.C, p, October

Chapter 6 : CiteSeerX â€" Autoepistemic Logic of Knowledge and Beliefs

An autoepistemic logic programming language is derived from a subset of a three-valued autoepistemic logic, called 3AEL. Autoepistemic programs generalize several ideas underlying logic programming: stable, supported, and *well-founded models, Fitting's semantics, Kunen's semantics, and abductive.*

Chapter 7 : Non-monotonic Logic (Stanford Encyclopedia of Philosophy)

We investigate the relationship between two epistemic nonmonotonic formalisms: autoepistemic logic and introspective circumscription. Finitely axiomatized autoepistemic theories are shown to be equivalent to the propositional case of introspective circumscription.

Chapter 8 : Robert C. Moore, Autoepistemic logic - PhilPapers

autoepistemic logic, and their basic properties known from the literature. Section 3 introduces the concept of stable theory, and relates the consequence operation induced by that concept to modal.