Lecture 2:
Kratzer, Poole, Conditionals and Means-End Reasoning

ESSLLI'16: Deontic modality: linguistic and logical perspectives on oughts and ends

Cleo Condoravdi & Leon van der Torre

Kratzer framework and conditionals

The formal framework introduced by Kratzer in 1981 is closely related to the work on DSDL in the sixties and seventies, such as the work of Hansson and Lewis discussed in the handout of Lecture 1. Her semantics is therefore sometimes referred to as the Lewis-Kratzer semantics. In the semantics of Kratzer, the preference ordering is not given, but it is constructed from two sets of propositions, called the modal base and the ordering source. The modal base gives the set of worlds of the preference ordering, namely the worlds satisfying all the sentences of the modal base. The ordering source gives the preference relation between the worlds, namely the subset ordering on the sentences of the ordering source. When the sentences of the ordering source are taken as obligations, this represents that a maximal number of obligations is fulfilled, or a minimal number of them is violated. Conditional obligations are then evaluated as in DSDL.

Poole system and abduction[1]

In the mid 80s it became clear that there are two different activities that are going on in nonmonotonic logic: explanation and prediction. Explanation is where the proposition to be explained (the explanandum) is an observation in the world, and we want to explain why this occurred (this is one formalization of what C.S.Peirce called abduction). The second, prediction is where the explanandum is unknown, and the problem is to determine whether to predict the explanandum or not (forming a formalization of default reasoning).

It became clear that these are quite different, but that they naturally fit together. A common reasoning strategy is, given an observation, to explain it (using normality and abnormality assumptions) and then make predictions (using normality assumptions). One of the early applications was in diagnosis. This was in the form of what is now called 'abductive diagnosis'.

More recent work on probabilistic Horn abduction has shown the close relationship between Bayesian conditioning with abduction followed by prediction. The probabilities in probabilistic Horn abduction give us a finer-grained control than just normality and abnormality assumptions, and also gives us a clean and simple model-theoretic semantics.

Definition of a Poole system[2]

Poole uses first order logic (FOL) as the underlying logic.

F is a set of closed formulae, which we are treating as "facts". We assume the facts are consistent. These are intended to be statements which are true in the intended interpretation. More precisely these are statements we are not prepared to give up for some application.

Δ is a set of formulae, called the set of possible hypotheses. Any ground instance of these can be used as a hypothesis if it is consistent.

Definition 1. *A scenario of F, Δ is a set $D \cup F$ where D is a set of ground instances of elements of Δ such that $D \cup F$ is consistent.*

Definition 2. *If g is a closed formula then an explanation of g from F, Δ is a scenario of F, Δ which implies g.*

That is, g is explainable from F, Δ if there is a set D of ground instances of elements of Δ such that

1. $F \cup D \models g$ and
2. $F \cup D$ is consistent

$F \cup D$ is an explanation of g. The first condition says that D predicts g, and the second says that D does not predict something which we know is false.

Definition 3. *An extension of F, Δ is the set of logical consequences of a maximal (with respect to set inclusion) scenario of F, Δ.*

The implementation of Poole's system, called Theorist, looks for minimal explanations.

Similarity between Kratzer and Poole

The similarity between Kratzer and Poole is striking, both take maximal consistent subsets to define their logics. This may not be as surprising at it is as first sight, because taking the subset ordering and thus maximal consistent subsets is a natural operation. For example, Rescher's logic of commands (1966) predates the development of DSDL, and thus

[1] This section is taken from Poole's website https://www.cs.ubc.ca/∼poole/theorist.html.

[2] This section is taken from Poole, A Logical Framework for Default Reasoning. 1984/1988.

also the work of Kratzer and Poole. Moreover, as we will see in Lecture 3, it was also used in the mid eighties to define the standard framework for theory change (or AGM belief revision).[3]

Example

Assume that to achieve goal proposition g, there are action propositions $do(a_1), do(a_2), \ldots, do(a_n)$. In other words, $do(a_i)$ logically implies g, and g implies the disjunction $do(a_1) \vee do(a_2) \vee \ldots \vee do(a_n)$.

Now, on the one hand, you want to maximize your goals, which is g. On the other hand, you want to do so with a minimal number of actions. That is, you do not want to derive $do(a_1) \wedge do(a_2) \wedge \ldots \wedge do(a_n)$. One way to model such means end reasoning, is to associate implicit costs with the actions, and there are implicit goals to minimize these costs.

Example 1. *Consider "If you want to go to Brooklyn, then you should take the A-train." In this example, g is "being in Brooklyn", a_1 is "taking the A train", a_2 is "walking to Brooklyn", etc. Moreover, let $c(a_1)$ be consequences of a_1 like paying money for ticket, $c(a_2)$ be consequences of a_2 like being very tired etc.*

$F = \{do(x) \rightarrow g, g \rightarrow do(a_1) \vee \ldots \vee do(a_n), do(x) \rightarrow c(x)\}$

$\Delta = \{\neg do(x), g\}$

$F \cup \Delta$ *implies a contradiction, so an explanation* $F \cup D$ *of* g *has to exclude one of the* $\neg do(a_i)$.

Likewise, you can create alternative orderings by representing the costs differently, for example:

$\Delta = \{\neg do(a_1), \neg do(a_1) \wedge \neg do(a_2), \ldots, g\}$

Makinson 1994/2005 sceptical approach[4]

The relation between Kratzer and Poole becomes more explicit in the work of Makinson and Freund. They refer to the 1994 version of Poole's system.[5]

As shown by Makinson in 1994[6] Poole's original liberal conception of the "extension family function" associated with the pair (A, K) can be modified to a sceptical approach, providing a preferential inference relation when the Poole system is one "without constraints", i.e. when its set K of constraints is empty. Such a Poole system can be identified with a set of prerequisite-free normal defaults in the sense of Reiter, and the associated preferential inference relation then corresponds to the sceptical Reiter extension of A. It was noticed by Makinson and Poole that the preferential inference relation associated with such a Poole system can be represented by a special kind of preferential model, where the set of states is the set of all worlds. Makinson mentioned that the converse of this property was not settled, and conjectured that it may hold.

Freund proved that it holds in 1998.

Multiple extension problem

The move from Poole's system to Makinson's non-monotonic logic is not as innocent as it may seem. An important property of a Poole system is that there may be multiple extensions. This idea is also present in many other non-monotonic logics, such as logic programming, Reiter's default logic and Dung's abstract argumentation. In answer set programming, it is even argued that the possibility of multiple extensions is one of the main advantages of such approaches (because there can be multiple solutions to constraint based problems, like multiple solutions to a Sudoku puzzle). However, traditionally the possibility of multiple extensions was seen as a problem, because we do not know which extension to choose.

In the example, we cannot derive yet that the best way to go to Brooklyn is by $do(a_1)$, because alternative explanations like walking there $do(a_2)$ exist as well. In general, this multiple extension problem can be handled in various ways:

1. We can add formulas to the facts to exclude some extensions. This solution has several drawbacks.
2. We can add constraints to the formal set up to exclude some extensions (see appendix). This is the way Poole himself handled it.
3. We can add preferences (priorities, or values) among Δ. This may be most natural from a decision theoretic perspective, but less natural for other applications.

Finally, there are some more sophisticated ways to achieve similar results. As mentioned above, in Theorist one can represent minimal explanations. An alternative approach to diagnosis (by Reiter, from first principles) would assume for each element of Δ that it is either true of false. This seems closer to Kratzer ordering. Finally, in the nineties the so-called causal approach became popular, which is based on the principle that every observation needs to be explained. This was implemented in the so-called causal calculator.

Ability

Boutilier[7] extends DSDL with a simple model of action between controllable and uncontrollable propositions, This makes it possible to distinguish between ought to be "there should be a fence" and ought to do "you should build a fence." We can define the ability of an agent as a possible action (or a sequence of actions in planning) such that a state is reached. Such models are used in decision and game theory, as well as in STIT logic. We can define that an agent

[3] More examples are given by David Makinson in "Five faces of minimality" (1993).

[4] This is taken from Freund, Preferential reasoning in the perspective of Poole default logic. AIJ, 1998.

[5] D. Poole, Default logic, in: Gabbay, Hogger and Robinson, eds., Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 3, Nonmonotonic and Uncertain Reasoning (Oxford University Press, Oxford, 1994).

[6] Section 3.31 of D. Makinson, General patterns in nonmonotonic reasoning, in: D. Gabbay, Hogger and Robinson, eds., Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. irl, Nonmonotonic and Uncertain Reasoning (Oxford University Press, Oxford, 1994) 35-l 10.

[7] Towards a logic of qualitative decision theory, KR 1994.

sees to it that he is in Brooklyn, or a set of agents sees to it that they are in Brooklyn.

In such a setting, it is tempting to define ought to do in terms of ought to be: you ough tto do something iff it ought to be that you do it. Horty (2001) shows that this does not work for STIT theory, and he defines an alternative so-called dominance ought instead.

Planning and non-monotonic logic

In our example, a single action could achieve the goal. We say that the goal is the effect or postcondition of the action. However, in most models of means-end reasoning, we can consider also sequences of actions, or plans. To define whether actions can be executed in sequence, we associate not only postconditions with actions, but also preconditions. The preconditions specify whether the action can be done.

In so-called classical planning, the algorithms find ways to order atomic actions in sequences such that on the one hand, the postcondition of each action imply the precondition of the next action, and on the other hand, after the last action, the goal is achieved. The prototypical example is the block worlds, where blocks can be put on top of each other to build stacks of blocks (mimicking the reasoning of a three year old).

In planning, we can reason to find a plan like $a_1; a_2; a_3$ in two ways. In forward reasoning or data driven approach, we first consider the actions we can do and choose a_1. Then we consider the actions we can do after a_1 and we choose a_2. Finally we consider the actions we can do after a_2 and we choose a_3. We reached the goal so we stop. In the backward reasoning approach, we consider which actions can achieve the goal, and we choose a_3, then we consider which actions can see to the preconditions of a_3, and we choose a_2. Similarly, we choose a_1. Only in the backward reasoning approach we use the goal in our search for a plan, and sometimes we refer to such backwards reasoning as means-end reasoning. Also combinations of forward and backward reasoning are used in classical planning.

Non-monotonic logic is used extensively to deal with action and change in planning. However, logic is not used for the combinatorial search of plans, but for the reasoning about the effects of actions. If the effect of taking the A train is that you are in Brooklyn, then we want to derive also that the rest of the world does not change after taking the A train. This is known as the frame problem.

An advantage of classical planning in terms of pre and postconditions is that it is similar to the model of states and state transitions is similar to the model of computing. So computer scientists like this model of action. However, a drawback of classical planning is that it is rather inefficient, as shown by planning competitions, so it is rarely used. Alternatives are based, for example, on action refinement, or graph planning. Classical planning mainly survives in models of folk psychology, for example in BDI theory.

If we abstract away from the search for plans, such as the choice between forward and backward reasoning, and we abstract away from the reasoning about the effects of plans, such as the frame problem, then we end up with a choice structure where at each moment the agents have to choose among their actions. If there are multiple agents then we have game theoretic structures. Instead of actions or plans, we have strategies, or conditional plans. They make it possible to analyse agent interaction at a higher level of abstraction, such as the prisoner's dilemma.

Appendix

Makinson's bridge logic: Assumption consequence

This section is based on Makinson 2003.[8]

Makinson presents Poole systems as a method to go from classical logic to non-monotonic logic, restricting the set of background assumptions (as alternatives: the set of valuations, or adding rules).

Pivotal-Assumption Consequence fixes a background set K of formulae. K behaves as the modal base in Kratzer's theory.

Definition 4. *$A \vdash_K x$ ($x \in Cn_K(A)$) iff $K \cup A \vdash x$.*

Class of all pivotal-assumption consequence relations: $\vdash_K$ for some set K

Theorem 1 (Representation)**.** *$\vdash_K$ is pivotal-assumption consequence iff it satisfies the following three properties (but not necessarily substitution!):*

1. *Paraclassical*
   - *Supraclassical (includes classical consequence)*
   - *Closure operation (reflexivity + idempotence + monotony)*
2. *Disjunction in premises (alias OR)*
3. *Compact*

Default-Assumption Consequence allows background assumptions K to vary with current premises A. Diminish K when inconsistent with A and work with maximal subsets of K that are consistent with A. K behaves as the ordering base in Kratzer's approach.

Definition 5. *$A|\!\sim_K x$ ($x \in C_K(A)$) iff $K' \cup A \vdash x$ for every subset $K' \subseteq K$ maxiconsistent with A.*

Freund's results

This is taken from Freund, Preferential reasoning in the perspective of Poole default logic. AIJ, 1998

The sceptical inference relation associated with a Poole system without constraints is known to have a simple semantic representation by means of a smooth order directly defined on the set of interpretations associated with the underlying language.

Conversely, Freund proves in his paper that, on a finite propositional language, any preferential inference relation defined by such a model is induced by a Poole system without constraints. In the particular case of rational relations, the associated set of defaults may be chosen to be minimal; it then consists of a set of formulae, totally ordered through

[8] Bridges between Classical and Nonmonotonic Logic. *Logic Journal of the IGPL* 11(1):69–96.

classical implication, with cardinality equal to the height of the given relation. This result can be applied to knowledge representation theory and corresponds, in revision theory, to Grove's family of spheres. In the framework of conditional knowledge bases and default extensions, it implies that any rational inference relation may be considered as the rational closure of a minimal knowledge base. An immediate consequence of this is the possibility of replacing any conditional knowledge base by a minimal one that provides the same amount of information.

Normative Systems

Normative systems can be developed as a generalisation of assumption and rule based consequence, where the input A is no longer necessarily in the output.

Definition 6. *$x \in out_K(A)$ iff $K' \vdash x$ for every subset $K' \subseteq K$ maxiconsistent with A.*

Theorem 2. $Cn_K(A) = Cn(A \cup out_K(A))$

Definitions Poole system with constraints

This section is taken from Poole, A Logical Framework for Default Reasoning. 1984/1988

Constraints are a way to say "this default should not be applicable in this case", without any side effect of doing this. They have been found to be a very useful mechanism in practice. Other systems overcome such derivations by allowing the defaults as rules which can only be used in one direction and have explicit exceptions [Reiter80] or by having fixed and variable predicates [McCarthy86]. The idea is to define a set of constraints used to prune the set of scenarios. They are just used to reject scenarios and cannot be used to explain anything. Constraints are formulae with which scenarios must be consistent.

We introduce a set C of closed formulae called the set of constraints. The definition of scenario is revised as follows:

Definition 7. *A scenario of F, C, Δ is a set $D \cup F$ where D is a set of ground instances of elements of Δ such that $D \cup F \cup C$ is consistent.*

Makinson's Rule consequence

Use additional rules.

Pivotal-Rule Consequence fixes a set R of rules, where a rule is any ordered pair (a, x) of formulae.

Definition 8. *$A \vdash_R x$ iff $x \in$ every superset of A closed under both Cn and R*

Class of all pivotal-rule consequence relations:
$\vdash_R$ for some set R of rules.

Theorem 3 (Representation)**.** *Pivotal-rule consequence iff following two properties (but not necessarily disjunction in premises):*

1. *Paraclassical*
2. *Compact*

Theorem 4. $\{$*pivotal assumption*$\}$ = $\{$*pivotal rule*$\} \cap \{$*OR*$\}$ = $\{$*pivotal rule*$\} \cap \{$*pivotal valuation*$\}$

Equivalent definitions of $Cn_R(A)$

- $\cap\{X \subseteq A : X = Cn(X) = R(X)\}$
- $\cup\{A_n | n < \omega\}$ with $A_1 = A$, $A_{n+1} = Cn(A_n \cup R(A_n))$
- $\cup\{A_n | n < \omega\}$ with $A_1 = A$, $A_{n+1} = Cn(A_n \cup \{x\})$ where (a, x) is first rule in $\langle R \rangle$ s.t. $a \in A_n$ but $x \notin A_n$ (in the case that there is no such rule: $A_{n+1} = Cn(A_n)$)

Default-Rule Consequence fixes an ordering $\langle R \rangle$ of R.

Definition 9. $C_{\langle R \rangle}(A) = \cup\{A_n : n < \omega\}$ *with*

- $A_1 = A$ *and*
- $A_{n+1} = Cn(A_n \cup \{x\})$

where (a, x) is first rule in $\langle R \rangle$ such that:
$a \in A_n, x \notin A_n$, and x is consistent with A_n
(if no such rule: $A_{n+1} = Cn(A_n)$)

Facts:

- The sets $C_{\langle R \rangle}(A)$ for an ordering $\langle R \rangle$ of R are precisely the Reiter extensions of A using the normal default rules (a, x) alias $(a; x/x)$.
- The ordering makes a difference.
- Standard inductive definition versus fixpoints.

Definition 10 (Sceptical operation)**.**
$C_R(A) = \cap\{C_{\langle R \rangle}(A) : \langle R \rangle$ *an ordering of* $R\}$

Many variants of default assumption, valuation and rule based consequence have been introduced.

Normative Systems

Normative systems can be developed as a generalisation of assumption and rule based consequence, where the input A is no longer necessarily in the output.

Method 1: Assumption output

Definition 11. *$x \in out_K(A)$ iff $K' \vdash x$ for every subset $K' \subseteq K$ maxiconsistent with A.*

Theorem 5. $Cn_K(A) = Cn(A \cup out_K(A))$

Method 2: Rule output

Definition 12.
$x \in out_R(A)$ iff $x \in R(B) = \{x \mid (a, x) \in R, a \in A\}$ for every superset B of A closed under both Cn and R

Class of all pivotal-rule consequence relations:
$out_R = \{(a, x) \mid x \in out_R(\{a\})\}$ for some set R of rules.

Theorem 6. $\{$*pivotal rule*$\}$ = $\{$*rule output*$\} \cap \{$*reflexivity*$\}$.

Theorem 7. $Cn_R(A) = Cn(A \cup out_R(A))$

Default-Rule output fixes an ordering $\langle R \rangle$ of R.

Definition 13. $C_{\langle R \rangle}(A) = \cup\{A_n : n < \omega\}$ *with*

- $A_1 = \emptyset$ *and*
- $A_{n+1} = Cn(A_n \cup \{x\})$

where (a, x) is first rule in $\langle R \rangle$ such that:
$a \in Cn(A \cup A_n), x \notin Cn(A \cup A_n)$, and x is consistent with $Cn(A \cup A_n)$
(if no such rule: $A_{n+1} = Cn(A_n)$)