Electronic Communications of the EASST
Volume 12 (2008)

Formal Modeling of Adaptive and Mobile Processes

On the Modelling of an Agent's Epistemic State and its Dynamic Changes

Christoph Beierle and Gabriele Kern-Isberner

17 pages

Guest Editors: Julia Padberg, Kathrin Hoffmann
Managing Editors: Tiziana Margaria, Julia Padberg, Gabriele Taentzer
ECEASST Home Page: http://www.easst.org/eceasst/ ISSN 1863-2122

On the Modelling of an Agent's Epistemic State and its Dynamic Changes

Christoph Beierle[1] **and Gabriele Kern-Isberner**[2]

[1] Fakultät für Mathematik und Informatik
FernUniversität in Hagen, Germany

[2] Fakultät für Informatik
Technische Universität Dortmund, Germany

Abstract: Given a set of unquantified conditionals considered as default rules or a set of quantified conditionals such as probabilistic rules, an agent can build up its internal epistemic state from such a knowledge base by inductive reasoning techniques. Besides certain (logical) knowledge, epistemic states are supposed to allow the representation of preferences, beliefs, assumptions etc. of an intelligent agent. If the agent lives in a dynamic environment, it has to adapt its epistemic state constantly to changes in the surrounding world in order to be able to react adequately to new demands. In this paper, we present a high-level specification of the CONDOR system that provides powerful methods and tools for managing knowledge represented by conditionals and the corresponding epistemic states of an agent. Thereby, we are able to elaborate and formalize crucial interdependencies between different aspects of knowledge representation, knowledge discovery, and belief revision. Moreover, this specification, using Gurevich's Abstract State Machines, provides the basis for a stepwise refinement development process of the CONDOR system based on the ASM methodology.

Keywords: knowledge representation, agent model, epistemic state, belief change, abstract state machine

1 Introduction

Any intelligent agent must be capable to represent knowledge internally and to reason about it. The internal epistemic state of an agent represents the full cognitive state of the agent, including its beliefs, assumptions and preferences. The epistemic state should allow to evaluate and to compare statements, formulas or possible worlds with respect to their plausibility, possibilty, necessity, probability, etc. Furthermore, any agent living in a dynamic environment must be able to dynamically change its own epistemic state in the light of new information. In this paper, we are concerned with formally modelling such an agent's epistemic state, how it can be initialzed from a given knowledge base and how it may evolve in time within a dynamically changing environment.

Commonsense and expert knowledge is most generally expressed by rules, connecting a precondition and a conclusion by an *if-then*-construction. If-then-rules are more formally denoted as conditionals, and often they occur in the form of probabilistic (quantitative) conditionals like

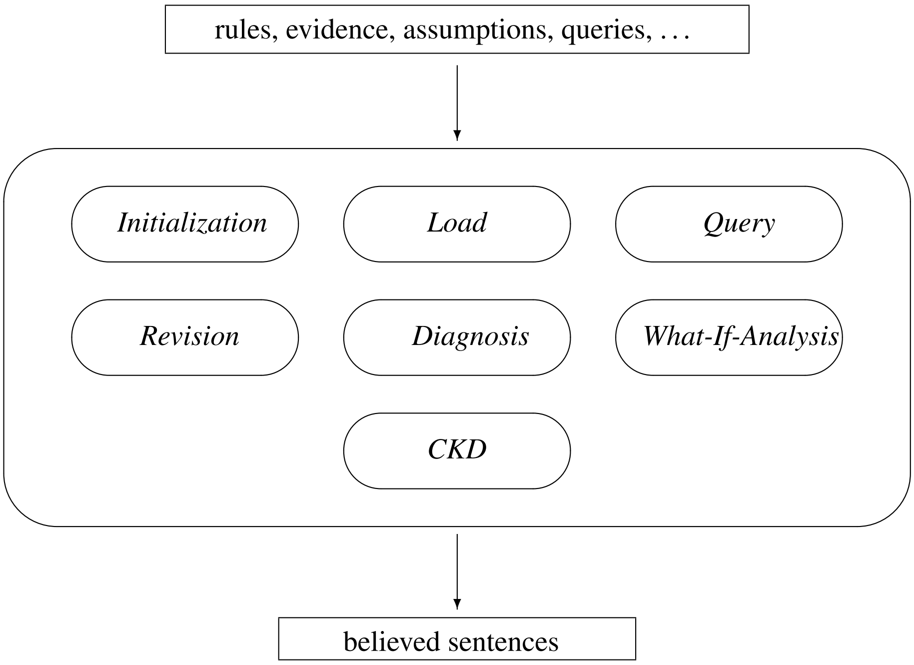

Figure 1: A bird's-eye view of the CONDOR-Systems

"*Students are young with a probability of (about) 80 %*" and "*Singles (i.e. unmarried people) are young with a probability of (about) 70 %*", where this commonsense knowledge can be expressed formally by $\{(young|student)[0.8], (young|single)[0.7]\}$. In another setting, qualitative conditionals like $(expensive|Mercedes)[n]$ are considered where $n \in \mathbb{N}$ indicates a degree of plausibility for the conditional "*Given that the car is a Mercedes, it is expensive*".

The crucial point with conditionals is that they carry generic knowledge which can be applied to different situations. This makes them most interesting objects in Artificial Intelligence, in theoretical as well as in practical respect. Within the CONDOR project (Conditionals - discovery and revision), we develop methods and tools for discovery and revision of knowledge expressed by conditionals.

Figure 1 provides a bird's-eye view of the CONDOR system. CONDOR can be seen as an agent being able to take rules, evidence, queries, etc., from the environment and giving back sentences it believes to be true with a degree of certainty. Basically, these degrees of belief are calculated from the agent's current epistemic state which is a representation of its cognitive state at the given time. The agent is supposed to live in a dynamic environment, so it has to adapt his epistemic state constantly to changes in the surrounding world and to react adequately to new demands.

In the following, we will develop a high-level, abstract model called CONDORASM, providing the top-level functionalities of the CONDOR system as they are indicated by the buttons in Figure 1. CONDORASM uses the methodology of Abstract State Machines (ASM) [Gur95, SSB01, BS03]. At the level of abstraction chosen for CONDORASM, we are not only able to decribe precisely the common functionalities for dealing with both quantitative and qualitative approaches. We also work out crucial interdependencies between e.g. inductive knowledge representation, knowledge discovery, and belief revision in a conditional setting. Moreover, CONDORASM provides the basis for a stepwise refinement development process of the CONDOR system.

The rest of this paper which is a revised and extended version of [BK03] is organized as follows: In Section 2, we provide very brief introductions to qualitative and quantitative logics and to the concept of Abstract State Machines. In Section 3, the universes of CONDORASM and its overall structure are introduced, while in Section 4 its top-level functions are specified. Section 5 contains some conclusions and points out further work.

2 Background

2.1 Qualitative and Quantitative Logic in a Nutshell

We start with a propositional language $\mathscr{L}$, generated by a finite set Σ of atoms $a,b,c,\ldots$. The formulas of $\mathscr{L}$ will be denoted by uppercase Roman letters $A,B,C,\ldots$. For conciseness of notation, we will omit the logical *and*-connector, writing AB instead of $A \wedge B$, and overlining formulas will indicate negation, i.e. $\overline{A}$ means $\neg A$. Let Ω denote the set of possible worlds over $\mathscr{L}$; Ω will be taken here simply as the set of all propositional interpretations over $\mathscr{L}$ and can be identified with the set of all complete conjunctions over Σ. The notation $\omega \models A$ means that the propositional formula $A \in \mathscr{L}$ holds in the possible world $\omega \in \Omega$.

By introducing a new binary operator $|$, we obtain the set $(\mathscr{L} \mid \mathscr{L}) = \{(B|A) \mid A,B \in \mathscr{L}\}$ of conditionals over $\mathscr{L}$. $(B|A)$ formalizes "*if A then B*" and establishes a plausible, probable, possible etc. connection between the *antecedent* A and the *consequent* B. Here, conditionals are supposed not to be nested, that is, antecedent and consequent of a conditional will be propositional formulas.

To give appropriate semantics to conditionals, they are usually considered within richer structures such as *epistemic states*. Besides certain (logical) knowledge, epistemic states also allow the representation of preferences, beliefs, assumptions etc of an intelligent agent. Basically, an epistemic state allows one to compare formulas or worlds with respect to plausibility, possibility, necessity, probability etc.

Well-known qualitative, ordinal approaches to represent epistemic states are Spohn's *ordinal conditional functions, OCFs*, (also called *ranking functions*) [Spo88], and *possibility distributions* [BDP92], assigning degrees of plausibility, or of possibility, respectively, to formulas and possible worlds. In such qualitative frameworks, a conditional $(B|A)$ is valid (or *accepted*), if its confirmation, AB, is more plausible, possible etc. than its refutation, $A\overline{B}$; a suitable degree of acceptance is calculated from the degrees associated with AB and $A\overline{B}$.

In a quantitative framework, most appreciated representations of epistemic states are provided by *probability functions* (or *probability distributions*) $P : \Omega \to [0,1]$ with $\sum_{\omega \in \Omega} P(\omega) = 1$. The probability of a formula $A \in \mathscr{L}$ is given by $P(A) = \sum_{\omega \models A} P(\omega)$, and the probability of a conditional $(B|A) \in (\mathscr{L} \mid \mathscr{L})$ with $P(A) > 0$ is defined as $P(B|A) = \dfrac{P(AB)}{P(A)}$, the corresponding conditional probability. Note that, since $\mathscr{L}$ is finitely generated, Ω is finite, too, and we only need additivity instead of σ-additivity.

2.2 Abstract State Machines

For a general introdcution to ASMs and also to stepwise refinement using ASMs we refer the reader to [Gur95], [Gur00], [SSB01], or [BS03]. Here we will only give a very brief summary, introducing the ASM concepts used in this paper and fixing our notation.

Each ASM has an ASM-signature Σ_{ASM} containing a finite collection of function names each of which has an arity. Nullary function names are also called constants. Function names are either *static* or *dynamic*, and each ASM signature contains the static constants *undef*, *true*, *false*.

A state $\mathscr{A}$ for an ASM signature Σ_{ASM} is a non-empty set X together with an interpretation $f^{\mathscr{A}} : X^n \rightarrow X$ for any n-ary function name f in Σ_{ASM}. $dom(f^A)$ is the set of tuples over X for which $f^{\mathscr{A}}$ yields a value different from *undef*. A function having only the values *true* and *false* is viewed as a relation.

The set X is also called *superuniverse*. The superuniverse may be devided into (sub-)universes which are represented by unary relations. The domain and codomain of a function name may also be given as a sequence of universes. For instance, if *EpState* is a universe (of epistemic states) and *currstate* is a nullary function name, the delaration

$$currstate : EpState$$

expresses that in any state $\mathscr{A}$ (the interpretation of) *currstate* yields an element of the universe *EpState*.

An *abstract state machine M* consists of an ASM signature Σ_{ASM}, an initial state for Σ_{ASM}, and a set of *transition rules*. For instance, $f(t_1, \ldots, t_n) := t_0$ with ground terms t_i is an *update rule*. Evaluating this rule in an ASM state $\mathscr{A}$ yields the (singleton) update set $U = \{(l, t_0^{\mathscr{A}}\}$ where the pair $l = (f, (t_1^{\mathscr{A}}, \ldots, t_n^{\mathscr{A}}))$ is a *location* in $\mathscr{A}$. Applying U to $\mathscr{A}$ results in a new state $\mathscr{A}'$ that coincides with $\mathscr{A}$ except that $f^{\mathscr{A}'}$ applied to the tuple $(t_1^{\mathscr{A}}, \ldots, t_n^{\mathscr{A}})$ yields $t_0^{\mathscr{A}}$. (Note that f must be a dynamic function; static functions are never changed.) For instance, in the state resulting from executing the update rule

$$currstate := inductive(rule\_base)$$

the nullary function *currstate* has a new value which is obtained by applying the function *inductive* to the nullary function *rule_base*.

There are also conditional rules, let rules, parallel executions rules, etc [BS03]. In general, a set of rules yields a set of updates. If this set is consistent – i.e., it does not contain different values for the same location – all updates are used to generate the new ASM state. A *run* of an ASM is a (finite or infinite) sequence of states, reflecting how the ASM evolves in time. It starts in its initial state, and each subsequent state is obtained from the previous one by applying all applicable rules simultaneously.

In the following sections, we will develop the abstract state machine CONDORASM which models an agent as sketched in Figure 1. The agent's interaction with the environment is modelled by using three kinds of dynamic functions (cf. [SSB01]):

controlled functions	:	changed by the ASM only
monitored functions	:	changed by the environment only
interaction functions	:	may be changed by both

Further details about the ASM methodology can be found in the references cited above.

3 The formal framework of CONDORASM

3.1 Universes

At a first and still very abstract level we do not distinguish between qualitative and quantitative conditionals. Therefore, we use $\mathcal{Q}$ as the universe of qualitative and quantitative scales.

The universe Σ of propositional variables provides a vocabulary for denoting simple facts. The universe Ω contains all possible worlds that can be distinguished using Σ. $\mathit{Fact}_{\mathcal{U}}$ is the set of all (unquantified) propositional sentences over Σ, i.e. $\mathit{Fact}_{\mathcal{U}}$ consists of all formulas from $\mathscr{L}$. The set of all (unquantified) conditional sentences from $(\mathscr{L} \mid \mathscr{L})$ is denoted by $\mathit{Rule}_{\mathcal{U}}$.

The universe of all sentences without any qualitative or quantitative measure is given by

$$\mathit{Sen}_{\mathcal{U}} = \mathit{Fact}_{\mathcal{U}} \cup \mathit{Rule}_{\mathcal{U}}$$

with elements written as A and $(B|A)$, respectively. Additionally, $\mathit{SimpleFact}_{\mathcal{U}}$ denotes the set of simple facts $\Sigma \subseteq \mathit{Fact}_{\mathcal{U}}$, i.e. $\mathit{SimpleFact}_{\mathcal{U}} = \Sigma$.

Analogously, in order to take quantifications of belief into account, we introduce the universe $\mathit{Sen}_{\mathcal{Q}}$ of all qualitative or quantitative sentences by setting

$$\mathit{Sen}_{\mathcal{Q}} = \mathit{Fact}_{\mathcal{Q}} \cup \mathit{Rule}_{\mathcal{Q}}$$

whose elements are written as $A[x]$ and $(B|A)[x]$, respectively, where $A, B \in \mathit{Fact}_{\mathcal{U}}$ and $x \in \mathcal{Q}$. For instance, the measured conditional $(B|A)[x]$ has the reading *if A then B with degree of belief x*. The set of measured simple facts is denoted by $\mathit{SimpleFact}_{\mathcal{Q}} \subseteq \mathit{Fact}_{\mathcal{Q}}$.

The universe of epistemic states is given by $\mathit{EpState} \subseteq \{\Psi \mid \Psi : \Omega \rightarrow \mathcal{Q}\}$. We assume that each $\Psi \in \mathit{EpState}$ uniquely determines a function (also denoted by Ψ) $\Psi : \mathit{Sen}_{\mathcal{U}} \rightarrow \mathcal{Q}$. For instance, in a probabilistic setting, for $\Psi = P : \Omega \rightarrow [0,1]$ we have $P(A) = \sum_{\omega \models A} P(\omega)$ for any unquantified sentence $A \in \mathit{Sen}_{\mathcal{U}}$.

Finally, there is a binary satisfaction relation (modelled by its characteristic function in the ASM framework) $\models_{\mathcal{Q}} \subseteq \mathit{EpState} \times \mathit{Sen}_{\mathcal{Q}}$ such that $\Psi \models_{\mathcal{Q}} S$ means that the state Ψ satisfies the sentence S. Typically, Ψ will satisfy a sentence like $A[x]$ if Ψ assigns to A the degree x (“*in Ψ, A has degree of probability / plausibility x*”).

In this paper, our standard examples for epistemic states are probability distributions, but note that the complete approach carries over directly to the ordinal framework (see eg. [KI01a]).

Example 1 *In a probabilistic setting, conditionals are interpreted via conditional probability. So for a probability distribution P, we have $P \models_{\mathcal{Q}} (B|A)[x]$ iff $P(B|A) = x$ (for $x \in [0,1]$).*

For instance, consider the three propositional variables a - being a student, b - being young, and c - being unmarried. *Let P be a probability distribution over the set of possible worlds Ω generated by $\Sigma = \{a,b,c\}$ with $P(abc) = 0.1950$, $P(ab\bar{c}) = 0.1758$, $P(a\bar{b}c) = 0.0408$, and $P(a\bar{b}\bar{c}) = 0.0519$. Then we have, e.g., $P \models_{\mathcal{Q}} (b|a)[0.8]$ since:*

$$P(b|a) = \frac{P(ab)}{P(a)} = \frac{0.1950 + 0.1758}{0.1950 + 0.1758 + 0.0408 + 0.0519} = \frac{0.3708}{0.4635} = 0.8$$

3.2 Overall structure

In the CONDORASM, the agent's current epistemic state is denoted by the controlled nullary function

$$currstate : EpState$$

The agents beliefs returned to the environment can be observed via the controlled function

$$believed\_sentences : \mathscr{P}(Sen_{\mathscr{Q}})$$

with $\mathscr{P}(S)$ denoting the power set of S.

As indicated in Figure 1, there are seven top-level functions that can be invoked, ranging from initialization of the system to the automatic discovery of conditional knowledge (CKD). Thus, we have a universe

$$\begin{aligned} WhatToDo \;=\; & \{Initialization, Load, Query, Revision, \\ & Diagnosis, What\text{-}If\text{-}Analysis, CKD\} \end{aligned}$$

The nullary interaction function

$$do : WhatToDo$$

is set by the environment in order to invoke a particular function. We tacitly assume that *do* is reset to *undef* by CONDORASM after each corresponding rule execution.

The appropriate inputs to the top-level functions are modelled by monitored nullary functions set by the environment. For instance, simply querying the system takes a set of (unquantified) sentences from $Sen_{\mathscr{U}}$, asking for the degree of belief for them. Similarly, the *What-If-Analysis* realizes hypothetical reasoning, taking a set of (quantified) sentences from $Sen_{\mathscr{Q}}$ as assumptions, together with a set of (unquantified) sentences from $Sen_{\mathscr{U}}$ as goals, asking for the degree of belief for these goals under the given assumptions. Figure 2 summarizes all monitored functions serving as inputs to the system; their specific usage will be explained in detail in the following section along with the corresponding top-level functionalities.

4 Top-level Functions in the CONDOR-System

4.1 Initialization

In the beginning, a prior epistemic state has to be built up on the basis of which the agent can start his computations. If no knowledge at all is at hand, simply the uniform epistemic state, modelled by the nullary function

$$uniform : EpState$$

is taken to initialize the system. For instance, in a probabilistic setting, this corresponds to the uniform distribution where all possible worlds have the same probability.

input type		monitored nullary function
$\mathscr{P}(Sen_{\mathscr{Q}})$	:	*rule_base* *new_information* *assumptions*
$\mathscr{P}(Sen_{\mathscr{U}})$	:	*queries* *goals*
$\mathscr{P}(Fact_{\mathscr{Q}})$	:	*evidence*
$\mathscr{P}(SimpleFact_{\mathscr{U}})$	:	*diagnoses*
EpState	:	*stored_state* *distribution*
RevisionOp	:	*rev_op*

Figure 2: Monitored function in CONDORASM

If, however, default knowledge or a set of probabilistic rules is at hand to describe the problem area under consideration, an epistemic state has to be found to appropriately represent this prior knowledge. To this end, we assume an inductive representation method to establish the desired connection between sets of sentences and epistemic states. Whereas generally, a set of sentences allows a (possibly large) set of models (or epistemic states), in an inductive formalism we have a function

$$inductive : \mathscr{P}(Sen_{\mathscr{Q}}) \rightarrow EpState$$

such that $inductive(S)$ selects a unique, "best" epistemic state from all those states satisfying S. Starting with the no-knowledge representing state *uniform* can be modelled by providing the empty set of rules since the constraint

$$uniform = inductive(\emptyset)$$

must hold.

Thus, we can initialize the system with an epistemic state by providing a set of (quantified) sentences S and generating a full epistemic state from it by inductively completing the knowledge given by S. For reading in such a set S, the monitored nullary function $rule\_base : \mathscr{P}(Sen_{\mathscr{Q}})$ is used:

if *do* = *Initialization*
  then *currstate* := *inductive*(*rule_base*)

Selecting a "best" epistemic state from all all those states satisfying a set of sentences S is an instance of a general problem which we call the *inductive representation problem* (cf. [BK05]). There are several well-known methods to model such an inductive formalism, a prominent one being the *maximum entropy* approach.

Example 2 *In a probabilistic framework, the* principle of maximum entropy *associates to a set $\mathscr{S}$ of probabilistic conditionals the unique distribution $P^* = MaxEnt(\mathscr{S})$ that satisfies all conditionals in $\mathscr{S}$ and has maximal entropy, i.e., $MaxEnt(\mathscr{S})$ is the (unique) solution to the maximization problem*

$$\arg\max H(P') = -\sum_{\omega} P'(\omega) \log P'(\omega) \tag{1}$$
$$\text{s.t. } P' \text{ is a probability distribution with } P' \models \mathscr{S}.$$

The rationale behind this is that $MaxEnt(\mathscr{S})$ represents the knowledge given by $\mathscr{S}$ most faithfully, i.e. without adding information unnecessarily (cf. [Par94, PV97, KI98]).

We will illustrate the maximum entropy method by a small example.

Example 3 *Consider again the three propositional variables a - being a student, b - being young, and c -* being unmarried *from Example 1. The commonsense knowledge "Students and unmarried people* are mostly *young" an agent may have can be expressed probabilistically e.g. by the set $\mathscr{S} = \{(b|a)[0.8], (b|c)[0.7]\}$ of conditionals. The MaxEnt-representation $P^* = MaxEnt(\mathscr{S})$ is given in the following table*[1]*:*

ω	$P^*(\omega)$	ω	$P^*(\omega)$	ω	$P^*(\omega)$	ω	$P^*(\omega)$
abc	0.1950	$ab\overline{c}$	0.1758	$a\overline{b}c$	0.0408	$a\overline{b}\overline{c}$	0.0519
$\overline{a}bc$	0.1528	$\overline{a}b\overline{c}$	0.1378	$\overline{a}\overline{b}c$	0.1081	$\overline{a}\overline{b}\overline{c}$	0.1378

4.2 Loading an epistemic state

Another way to initialize the system with an epistemic state is to load such a state directly from the environment (where it might have been stored during a previous run of the system; this could be modelled easily by an additional top-level function). Therefore, there is a monitored nullary function *stored_state* : *EpState* which is used in the following rule:

if *do* = *Load*
  then *currstate* := *stored_state*

4.3 Querying an epistemic state

The function

$$belief : EpState \times \mathscr{P}(Sen_{\mathscr{U}}) \rightarrow \mathscr{P}(Sen_{\mathscr{Q}})$$

is the so-called belief measure function which is subject to the condition

$$belief(\Psi, \mathscr{S}) = \{S[x] \mid S \in \mathscr{S} \text{ and } \Psi \models_{\mathscr{Q}} S[x]\}$$

for every $\Psi \in EpState$ and $S \subseteq Sen_{\mathscr{U}}$. For a given state Ψ, the call $belief(\Psi, S)$ returns, in the form of measured sentences, the beliefs that hold with regard to the set of basic sentences $S \subseteq$

[1] $MaxEnt(\mathscr{S})$ has been computed with the expert system shell SPIRIT [RK97a, RK97b, RRK06].

$Sen_{\mathscr{U}}$. The monitored function $queries : \mathscr{P}(Sen_{\mathscr{U}})$ holds the set of sentences and is used in the rule:

if *do* = *Query*
then $believed\_sentences := belief(currstate, queries)$

Example 4 *Suppose the current epistemic state is* $currstate = MaxEnt(\mathscr{S})$ *from Example* 3 *above, and our query is* "What is the probability that unmarried students are young?", *i.e.* $queries = \{(b|ac)\}$. *The system returns* $belief(currstate, queries) = \{(b|ac)[0.8270]\}$, *that is, unmarried students are supposed to be young with probability 0.8270.*

4.4 Revision of Conditional Knowledge

Belief revision, the theory of dynamics of knowledge, has been mainly concerned with propositional beliefs for a long time. The most basic approach here is the *AGM-theory* presented in the seminal paper [AGM85] as a set of postulates outlining appropriate revision mechanisms in a propositional logical environment. This framework has been widened by Darwiche and Pearl [DP97a] for (qualitative) epistemic states and conditional beliefs. An even more general approach, unifying revision methods for quantitative and qualitative representations of epistemic states, is described in [KI01a]. The crucial meaning of conditionals as *revision policies* for belief revision processes is made clear by the so-called *Ramsey test* [Ram50], according to which a conditional $(B|A)$ is accepted in an epistemic state Ψ, iff revising Ψ by A yields belief in B:

$$\Psi \models (B|A) \quad \text{iff} \quad \Psi * A \models B \tag{2}$$

where $*$ is a belief revision operator (see e.g. [Ram50, BG93]).

Note, that the term "belief revision" is a bit ambiguous: On the one hand, it is used to denote quite generally *any* process of changing beliefs due to incoming new information [Gär88]. On a more sophisticated level, however, one distinguishes between different kinds of belief change. Here, *(genuine) revision* takes place when new information about a static world arrives, whereas *updating* tries to incorporate new information about a (possibly) evolving, changing world [KM91]. *Expansion* simply adds new knowledge to the current beliefs, in case that there are no conflicts between prior and new knowledge [Gär88]. *Focusing* [DP97b] means applying generic knowledge to the evidence present by choosing an appropriate context or reference class. *Contraction* [Gär88] and *erasure* [KM91] are operations inverse to revision and updating, respectively, and deal with the problem of how to "forget" knowledge. In this paper, we will make use of this richness of different operations, but only on a surface level, without going into details. The explanations given above will be enough for understanding the approach to be developed here. An interested reader may follow the mentioned references. For a more general approach to belief revision both in a symbolic and numerical framework, cf. [KI01a]. The revision operator $*$ used above is most properly looked upon as a *revision or updating* operator. We will stick, however, to the term *revision*, and will use it in its general meaning, if not explicitly stated otherwise.

The universe of revision operators is given by

$$RevisionOp = \{Update, Revision, Expansion, Contraction, Erasure, Focusing\}$$

and the general task of revising knowledge is realized by a function

$$revise : EpState \times RevisionOp \times \mathscr{P}(Sen_{\mathscr{Q}}) \rightarrow EpState$$

A call $revise(\Psi, op, \mathscr{S})$ yields a new state where Ψ is modified according to the revision operator op and the set of sentences $\mathscr{S}$. Note that we consider here belief revision in a very general and advanced form: We revise epistemic states by sets of conditionals – this exceeds the classical AGM-theory by far which only deals with sets of propositional beliefs.

The constraints the function *revise* is expected to satisfy depend crucially on the kind of revision operator used in it and also on the chosen framework (ordinal or e.g. probabilistic). Therefore, we will merely state quite basic constraints here, which are in accordance with the AGM theory [AGM85] and its generalizations [DP97a, KI01a].

The first and most basic constraint corresponds to the *success postulate* in belief revision theory: if the change operator is one of $Update, Revision, Expansion$, the new information is expected to be present in the posterior epistemic state:

$$\begin{aligned} revise(\Psi, Revision, \mathscr{S}) &\models_{\mathscr{Q}} \mathscr{S} \\ revise(\Psi, Update, \mathscr{S}) &\models_{\mathscr{Q}} \mathscr{S} \\ revise(\Psi, Expansion, \mathscr{S}) &\models_{\mathscr{Q}} \mathscr{S} \end{aligned}$$

Furthermore, any revision process should satisfy *stability* – if the new information to be incorporated is already represented in the present epistemic state, then no change shall be made:

$$\text{If } \Psi \models_{\mathscr{Q}} \mathscr{S} \text{ then:}$$

$$\begin{aligned} revise(\Psi, Revision, \mathscr{S}) &= \Psi \\ revise(\Psi, Update, \mathscr{S}) &= \Psi \\ revise(\Psi, Expansion, \mathscr{S}) &= \Psi \end{aligned}$$

Similarly, for the deletion of information we get:

$$\text{If not } \Psi \models_{\mathscr{Q}} \mathscr{S} \text{ then:}$$

$$\begin{aligned} revise(\Psi, Contraction, \mathscr{S}) &= \Psi \\ revise(\Psi, Erasure, \mathscr{S}) &= \Psi \end{aligned}$$

To establish a connection between revising and retracting operations, one may further impose *recovery constraints*:

$$\begin{aligned} revise(revise(\Psi, Contraction, \mathscr{S}), Revision, \mathscr{S}) = \Psi \\ revise(revise(\Psi, Erasure, \mathscr{S}), Update, \mathscr{S}) = \Psi \end{aligned}$$

A correspondence between inductive knowledge representation and belief revison can be established by the condition

$$inductive(S) = revise(uniform, Update, S). \tag{3}$$

Thus, inductively completing the knowledge given by S can be taken as revising the non-knowledge representing epistemic state *uniform* by updating it to S.

In CONDORASM, revision is realized by the rule

if *do* = *Revision*
  then *currstate* := *revise*(*currstate*, *rev_op*, *new_information*)

where the monitored functions $rev\_op : RevisionOp$ and $new\_information : \mathscr{P}(Sen_{\mathscr{Q}})$ provide the type of revision operator to be applied and the set of new sentences to be taken into account, respectively.

Example 5 *In a probabilistic framework, a powerful tool to revise (more appropriately: update) probability distributions by sets of probabilistic conditionals is provided by the* principle of minimum cross-entropy *which generalizes the principle of maximum entropy in the sense of (3): Given a (prior) distribution, P, and a set, $\mathscr{S}$, of probabilistic conditionals, The MinEnt-*distribution $P_{ME} = MinEnt(P, \mathscr{S})$ *is the (unique) distribution that satisfies all constraints in $\mathscr{S}$ and has minimal cross-entropy with respect to P, i.e. P_{ME} solves the minimization problem*

$$\arg\min R(P', P) = \sum_{\omega} P'(\omega) \log \frac{P'(\omega)}{P(\omega)} \tag{4}$$
$$\text{s.t. } P' \text{ is a probability distribution with } P' \models \mathscr{S}$$

*If $\mathscr{S}$ is basically compatible with P (i.e. P-*consistent*, cf. [KI01a]), then P_{ME} is guaranteed to exist (for further information and lots of examples, see [Csi75, PV92, Par94, KI01a]). The cross-entropy between two distributions can be taken as a directed (i.e. asymmetric) information distance [Sho86] between these two distributions. So, following the principle of minimum cross-entropy means to revise the prior epistemic state P in such a way as to obtain a new distribution which satisfies all conditionals in $\mathscr{S}$ and is as close to P as possible. So, the MinEnt-principle yields a probabilistic belief revision operator, $*_{ME}$, associating to each probability distribution P and each P-consistent set $\mathscr{S}$ of probabilistic conditionals a revised distribution $P_{ME} = P *_{ME} \mathscr{S}$ in which $\mathscr{S}$ holds.*

Example 6 *Suppose that some months later, the agent from Example 3 has changed his mind concerning his formerly held conditional belief $(young|student)$ – he now believes that students are young with a probability of 0.9. So an updating operation has to modify P^* appropriately. We use MinEnt-revision to compute $P^{**} = revise(P^*, Update, \{(b|a)[0.9]\})$. The result is shown in the table below.*

ω	$P^{**}(\omega)$	ω	$P^{**}(\omega)$	ω	$P^{**}(\omega)$	ω	$P^{**}(\omega)$
abc	0.2151	$ab\bar{c}$	0.1939	$a\bar{b}c$	0.0200	$a\bar{b}\bar{c}$	0.0255
$\bar{a}bc$	0.1554	$\bar{a}b\bar{c}$	0.1401	$\bar{a}\bar{b}c$	0.1099	$\bar{a}\bar{b}\bar{c}$	0.1401

It is easily checked that indeed, $P^{**}(b|a) = 0.9$ (only approximately, due to rounding errors, since

$$P^{**}(b|a) = \frac{P^{**}(ab)}{P^{**}(a)} = \frac{0.2151 + 0.1939}{0.2151 + 0.1939 + 0.0200 + 0.0255} = \frac{0.4090}{0.4545} \approx 0.8999$$

holds).

4.5 Diagnosis

Having introduced these first abstract functions for belief revision, we are already able to introduce additional functions. As an illustration, consider the function

$$diagnose : EpState \times \mathscr{P}(Fact_{\mathscr{Q}}) \times \mathscr{P}(SimpleFact_{\mathscr{U}}) \to \mathscr{P}(SimpleFact_{\mathscr{Q}})$$

asking about the status of certain simple facts $D \subseteq SimpleFact_{\mathscr{U}} = \Sigma$ in a state Ψ under the condition that some particular factual knowledge $\mathscr{S}$ (so-called evidential knowledge) is given. It is defined by

$$diagnose(\Psi, \mathscr{S}, D) = belief(revise(\Psi, Focusing, \mathscr{S}), D)$$

Thus, making a diagnosis in the light of some given evidence corresponds to what is believed in the state obtained by adapting the current state by focusing on the given evidence.

Diagnosis is realized by the rule

if *do* = *Diagnosis*
then *believed_sentences* := *diagnose*(*currstate*, *evidence*, *diagnoses*)

where the monitored functions $evidence : \mathscr{P}(Fact_{\mathscr{Q}})$ and $diagnoses : \mathscr{P}(SimpleFact_{\mathscr{U}})$ provide the factual evidence and a set of (unquantified) facts for which a degree of belief is to be determined.

Example 7 *In a probabilistic framework, focusing on a certain evidence is usually done by conditioning the present probability distribution correspondingly. For instance, if there is certain evidence for being a student and being unmarried – i.e.* $evidence = \{student \wedge unmarried[1]\}$ *– and we ask for the degree of belief of being young – i.e.* $diagnoses = \{young\}$ *– for* $currstate = P^*$ *from Example 3, the system computes*

$$diagnose(P^*, \{student \wedge unmarried[1]\}, \{young\}) = \{young[0.8270]\}$$

and update believed_sentences to this set. Thus, if there is certain evidence for being an unmarried student, then the degree of belief for being young is 0.8270.

4.6 What-If-Analysis: Hypothetical Reasoning

There is a close relationship between belief revision and generalized nonmonotonic reasoning described by

$$\mathscr{R} \mid\!\sim_{\Psi} \mathscr{S} \quad \text{iff} \quad \Psi * \mathscr{R} \models_{\mathscr{Q}} \mathscr{S}$$

(cf. [KI01a]). In this formula, the operator $*$ may be a revision or an update operator. Here, we will use updating as the operation to study default consequences. So, *hypothetical reasoning* carried out by the function

$$nmr_{upd} : EpState \times \mathscr{P}(Sen_{\mathscr{Q}}) \times \mathscr{P}(Sen_{\mathscr{U}}) \to \mathscr{P}(Sen_{\mathscr{Q}})$$

can be defined by combining the *belief*-function and the *revise*-function:

$$nmr_{upd}(\Psi, \mathscr{S}_q, \mathscr{S}_u) \;=\; belief(revise(\Psi, Update, \mathscr{S}_q), \mathscr{S}_u)$$

The assumptions $\mathscr{S}_q$ used for hypotherical reasoning are being hold in the monitored function $assumptions : \mathscr{P}(Sen_{\mathscr{Q}})$ and the sentences $\mathscr{S}_u$ used as goals for which we ask for the degree of belief are being hold in the monitored function $goals : \mathscr{P}(Sen_{\mathscr{U}})$. Thus, we obtain the rule

if *do* = *What-If-Analysis*
then $believed\_sentences := nmr_{upd}(currstate, assumptions, goals)$

Example 8 *With this function, hypothetical reasoning can be done as is illustrated e.g. by "Given* P^* *in Example 3 as present epistemic state – i.e.* $currstate = P^*$ *–, what would be the probability of* $(b|c)$ *– i.e.* $goals = \{(b|c)\}$ *–, provided that the probability of* $(b|a)$ *changed to* 0.9 *– i.e.* $assumptions = \{(b|a)[0.9]\}$ *?"* CONDOR*'s answer is* $believed\_sentences = \{(b|c)[0.7404]\}$ *which corresponds to the probability given by* P^{**} *from Example 6.*

4.7 Conditional Knowledge Discovery

Conditional knowledge discovery is modelled by a function

$$CKD : EpState \rightarrow \mathscr{P}(Sen_{\mathscr{Q}})$$

that extracts measured facts and rules from a given state Ψ that hold in that state, i.e. $\Psi \models_{\mathscr{Q}} CKD(\Psi)$. More significantly, $CKD(\Psi)$ should be a set of "interesting" facts and rules, satisfying e.g. some minimality requirement. In the ideal case, $CKD(\Psi)$ yields a set of measured sentences that has Ψ as its "designated" representation via an inductive representation formalism *inductive*. Therefore, discovering most relevant relationships in the formal representation of an epistemic state may be taken as solving the *inverse representation problem* (cf. [KI00, BK02, FKB07]):

> Given an epistemic state Ψ find a set of (relevant) sentences S that has Ψ as its designated representation, i.e. such that $inductive(S) = \Psi$.

The intended relationship between the two operations *inductive* and *CKD* can be formalized by the condition

$$inductive(CKD(\Psi)) \;=\; \Psi$$

which holds for all epistemic states Ψ.

This is the theoretical basis for our approach to knowledge discovery. In practice, however, usually the objects knowledge discovery techniques deal with are not epistemic states but statistical data. We presuppose here that these data are at hand as some kind of distribution, e.g. as a frequency distribution or an ordinal distribution (for an approach to obtain possibility distributions from data, cf. [GK97]). These distributions will be of the same type as epistemic states in the corresponding framework, but since they are ontologically different, we prefer to introduce another term for the arguments of *CKD*-functions:

$$distribution : EpState$$

if *do* = *CKD*
then *believed_sentences* := *CKD*(*distribution*)

For instance, in a quantitative setting, Ψ may be a (rather complex) full probabilistic distribution over a large set of propositional variables. On the other hand, $CKD(\Psi)$ should be a (relatively small) set of probabilistic facts and conditionals that can be used as a faithful representation of the relevant relationships inherent to Ψ, e.g. with respect to the *MaxEnt*-formalism (cf. Section 4.1). So, the inverse representation problem for $inductive = MaxEnt$ reads like this: Find a set $CKD(\Psi)$ such that Ψ is the uniquely determined probability distribution satisfying $\Psi = MaxEnt(CKD(\Psi))$.

We will illustrate the basic idea of how to solve this *inverse MaxEnt-problem* by continuing Example 3.

Example 9 *The probability distribution we are going to investigate is P^* from Example 3. Starting with observing relationships between probabilities like*

$$P^*(\overline{a}b\overline{c}) = P^*(\overline{a}\overline{b}\overline{c}),$$

$$\frac{P^*(abc)}{P^*(ab\overline{c})} = \frac{P^*(\overline{a}bc)}{P^*(\overline{a}b\overline{c})},$$

$$\frac{P^*(a\overline{b}c)}{P^*(a\overline{b}\overline{c})} = \frac{P^*(\overline{a}\overline{b}c)}{P^*(\overline{a}\overline{b}\overline{c})},$$

the procedure described in [KI01b] yields the set $\mathscr{S}_u = \{(b|a),(b|c)\}$ of unquantified (structural) conditionals not yet having assigned any probabilities to them. Associating the proper probabilities (which are directly computable from P^*) with these structural conditionals, we obtain

$$\mathscr{S} \quad = \quad CKD(P^*) \quad = \quad \{(b|a)[0.8],\ (b|c)[0.7]\}$$

as a MaxEnt-generating set for P^*, i.e. $P^* = MaxEnt(\mathscr{S})$. In other words, the probabilistic conditionals

(young|student)[0.8],
(young|unmarried)[0.7]

that have been generated from P^* fully automatically, constitute a consise set of uncertain rules that faithfully represent the complete distribution P^* in an information-theoretically optimal way. So indeed, we arrived at the same set of conditionals we used to build up P^* in Example 3. Thus, in this case we have

$$CKD(MaxEnt(\mathscr{S})) = \mathscr{S}$$

But note, that in general, $CKD(P)$ will also contain redundant rules so that only

$$CKD(MaxEnt(\mathscr{S})) \supseteq \mathscr{S}$$

will hold.

5 Conclusions and Further Work

Starting from a bird's-eye view of the CONDOR system, we developed a high-level ASM specification for a system that provides powerful methods and tools for the modelling of an agent's epistemic state and its dynamic changes. Thereby, we were able to elaborate crucial interdependencies between different aspects of knowledge representation, knowledge discovery, and belief revision.

Whereas in this paper, we deliberately left the universe $\mathcal{Q}$ of quantitative and qualitative scales abstract, aiming at a broad applicability of our approach, in a further development step described in [BK07], we refined CONDORASM to a qualitative approach using unquantified conditionals as default rules and using ordinal conditional functions as its models [Spo88]. In vertical refinement steps, the abstract functions like *belief* and *revise* are elaborated by realizing them at lower-level data structures, following the ASM idea of stepwise refinement.

Acknowledgements: The research reported here was supported by the DFG – Deutsche Forschungsgemeinschaft (grants BE 1700/5-1, BE 1700/7-1 and KE 1413/2-1).

Bibliography

[AGM85] C. Alchourrón, P. Gärdenfors, P. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic* 50(2):510–530, 1985.

[BDP92] S. Benferhat, D. Dubois, H. Prade. Representing default rules in possibilistic logic. In *Proceedings 3th International Conference on Principles of Knowledge Representation and Reasoning KR'92*. Pp. 673–684. 1992.

[BG93] C. Boutilier, M. Goldszmidt. Revision by conditional beliefs. In *Proceedings 11th National Conference on Artificial Intelligence (AAAI'93)*. Pp. 649–654. Washington, DC., 1993.

[BK02] C. Beierle, G. Kern-Isberner. Knowledge discovery and the inverse representation problem. In *Proceedings of the 2002 International Conference on Information and Knowledge Engineering, IKE'02*. 2002.

[BK03] C. Beierle, G. Kern-Isberner. Modelling Conditional Knowledge Discovery and Belief Revision by Abstract State Machines. In Börger et al. (eds.), *International Workshop on Abstract State Machines - ASM2003*. Lecture Notes in Computer Science 2589, pp. 186–2003. Springer-Verlag, 2003.

[BK05] C. Beierle, G. Kern-Isberner. Footprints of Conditionals. In Hutter and Stephan (eds.), *Mechanizing Mathematical Reasoning*. Lecture Notes in Computer Science 2605, pp. 99–119. Springer-Verlag, 2005.

[BK07] C. Beierle, G. Kern-Isberner. An ASM Refinement and Implementation of the Condor System using Ordinal Conditional Functions. In Prinz (ed.), *Proceedings 14th International Workshop on Abstract State Machines (ASM'2007)*. Agder University College, Grimstad, Norway, 2007.

[BS03] E. Börger, R. Stärk. *Abstract State Machines: A Method for High-Level System Design and Analysis*. Springer-Verlag, 2003.

[Csi75] I. Csiszár. I-divergence geometry of probability distributions and minimization problems. *Ann. Prob.* 3:146–158, 1975.

[DP97a] A. Darwiche, J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence* 89:1–29, 1997.

[DP97b] D. Dubois, H. Prade. Focusing vs. Belief Revision: A fundamental distinction when dealing with generic knowledge. In *Proceedings First International Joint Conference on Qualitative and Quantitative Practical Reasoning, ECSQARU-FAPR'97*. Pp. 96–107. Springer, Berlin Heidelberg New York, 1997.

[FKB07] J. Fisseler, G. Kern-Isberner, C. Beierle. Learning Uncertain Rules with CONDOR-CKD. In *Proceedings 20th International FLAIRS Conference, FLAIRS'07*. Pp. 74–79. AAAI Press, Menlo Park, California, 2007.

[Gär88] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge, Mass., 1988.

[GK97] J. Gebhardt, R. Kruse. Background and perspectives of possibilistic graphical models. In *Proceedings First International Joint Conference on Qualitative and Quantitative Practical Reasoning, ECSQARU-FAPR'97,*. Pp. 108–121. Springer, 1997.

[Gur95] Y. Gurevich. Evolving Algebras 1993: Lipari Guide. In Börger (ed.), *Specification and Validation Methods*. Pp. 9–36. Oxford University Press, 1995.

[Gur00] Y. Gurevich. Sequential abstract-state machines capture sequential algorithms. *ACM Trans. Comput. Log.* 1(1):77–111, 2000.

[KI98] G. Kern-Isberner. Characterizing the principle of minimum cross-entropy within a conditional-logical framework. *Artificial Intelligence* 98:169–208, 1998.

[KI00] G. Kern-Isberner. Solving the inverse representation problem. In *Proceedings 14th European Conference on Artificial Intelligence, ECAI'2000*. Pp. 581–585. IOS Press, Berlin, 2000.

[KI01a] G. Kern-Isberner. *Conditionals in nonmonotonic reasoning and belief revision*. Springer, Lecture Notes in Artificial Intelligence LNAI 2087, 2001.

[KI01b] G. Kern-Isberner. Discovering most informative rules from data. In *Proceedings International Conference on Intelligent Agents, Web Technologies and Internet Commerce, IAWTIC'2001*. 2001.

[KM91] H. Katsuno, A. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proceedings Second International Conference on Principles of Knowledge Representation and Reasoning, KR'91*. Pp. 387–394. Morgan Kaufmann, San Mateo, Ca., 1991.

[Par94] J. Paris. *The uncertain reasoner's companion – A mathematical perspective*. Cambridge University Press, 1994.

[PV92] J. Paris, A. Vencovská. A method for updating that justifies minimum cross entropy. *International Journal of Approximate Reasoning* 7:1–18, 1992.

[PV97] J. Paris, A. Vencovska. In defence of the maximum entropy inference process. *International Journal of Approximate Reasoning* 17(1):77–103, 1997.

[Ram50] F. Ramsey. General propositions and causality. In Braithwaite (ed.), *Foundations of Mathematics and other logical essays*. Pp. 237–257. Routledge and Kegan Paul, New York, 1950.

[RK97a] W. Rödder, G. Kern-Isberner. Léa Sombé und entropie-optimale Informationsverarbeitung mit der Expertensystem-Shell SPIRIT. *OR Spektrum* 19(3):41–46, 1997.

[RK97b] W. Rödder, G. Kern-Isberner. Representation and extraction of information by probabilistic logic. *Information Systems* 21(8):637–652, 1997.

[RRK06] W. Rödder, E. Reucher, F. Kulmann. Features of the Expert-System-Shell SPIRIT. *Logic Journal of the IGPL* 14(3):483–500, 2006.

[Sho86] J. Shore. Relative entropy, probabilistic inference and AI. In Kanal and Lemmer (eds.), *Uncertainty in Artificial Intelligence*. Pp. 211–215. North-Holland, Amsterdam, 1986.

[Spo88] W. Spohn. Ordinal conditional functions: a dynamic theory of epistemic states. In Harper and Skyrms (eds.), *Causation in Decision, Belief Change, and Statistics, II*. Pp. 105–134. Kluwer Academic Publishers, 1988.

[SSB01] R. Stärk, J. Schmid, E. Börger. *Java and the Java Virtual Machine: Definition, Verification, Validation*. Springer-Verlag, 2001.