From knowledge-based programs to graded belief-based programs – Part I: On-line reasoning

Noël Laverny and Jérôme Lang
IRIT
Université Paul Sabatier
31062 Toulouse Cedex, France
laverny@irit.fr, lang@irit.fr

Abstract. Knowledge-based programs ([10, 18]) are a powerful notion for expressing action policies in which branching conditions refer to implicit knowledge. However, branching conditions in knowledge-based programs cannot refer to possibly erroneous beliefs or to graded belief, such as "if my belief that φ holds is high then do some action α else perform some sensing action β". The purpose of this paper is to build a framework where such programs can be expressed. In this paper we focus on the execution of such a program (a companion paper investigates issues relevant to the off-line evaluation and construction of such programs). We define a simple graded version of doxastic logic KD45 as the basis for the definition of belief-based programs. Then we study the way the agent's belief state is maintained when executing such programs, which calls for revising belief states by observations (possibly unreliable or imprecise) and progressing belief states by physical actions (which may have normal as well as exceptional effects).

1 Introduction

Knowledge-based programs, or KBPs ([10] – see also [18, 2, 21, 14, 3, 15]), are a powerful notion for expressing action policies in which branching conditions refer to implicit knowledge and call for a deliberation task at execution time. However, branching conditions in KBPs cannot refer to possibly erroneous beliefs or to graded belief, such as "`while` I have no strong belief about the direction of the railway station `do` ask someone". The purpose of this paper is to build a framework for such *belief-based programs* (BBPs).

While knowledge states in KPBs are expressed in epistemic logic (usually S5), BBPs need a logic of *graded belief*, where different levels of uncertainty or entrenchment can be expressed. We therefore have to commit to a choice regarding the nature of uncertainty we wish to handle. Rather than reasoning with probabilistic belief states (and therefore introducing probabilistic modalities), which would take us far from usual logics of knowledge or belief such as S5 and KD45, we choose to define belief states as *ordinal conditional functions* (OCF) [20] – also called *kappa-fonctions*. Introducing OCFs in logic is easy (see [12, 5, 6] for logical frameworks of dynamicity and uncertainty based on OCFs); besides, OCFs are expressive enough in many situations where there exists only a small number of "belief degrees"; therefore they are a good trade-off between simplicity and expressivity, as well as between ordinality and cardinality, since they allow for an approximation of probabilities without the technical difficulties raised by the integration of logic and probability.

A graded version of KD45 is defined in Section 2. In Section 3 we show how belief states are *revised* by possibly unreliable observations produced by sensing actions. In Section 4 we show how belief states are *progressed* when the agent performs *(physical) actions* which may have alternative effects, some of which being more exceptional than others. Belief-based programs are defined in Section 5. Section 6 discusses related work.

2 KD45$_G$

The language $\mathcal{L}$ of graded doxastic dynamic logic KD45$_G$ is built from a *finite* set of propositional symbols, the usual connectives, the symbols $\top$ et $\bot$, and the doxastic modalities $\mathbf{B_1}$, $\mathbf{B_2}$, …, $\mathbf{B}_\infty$. A formula of $\mathcal{L}$ is *flat* if it does not contain nested modalities and *objective* if it does not contain any modality. *For the purposes of this paper it is sufficient to focus on flat formulas only* (non-flat formulas do not play any role for reasoning when executing belief-based programs).

Formulas of KD45$_G$ are denoted by capital Greek letters Φ, Ψ etc. while objective formulas are denoted by small Greek letters φ, ψ etc. $\mathbf{B_i}\varphi$ intuitively means that the agent belives φ with strength i. The larger i, the stronger the belief expressed by $\mathbf{B_i}$, and $\mathbf{B}_\infty$ is a *knowledge* modality and will be denoted more simply by $\mathbf{K}$ (belief with infinite strength is true knowledge).

Let us now define specific classes of flat formulas. A *doxastic atom* is a formula $\mathbf{B_i}\varphi$ where φ is objective. A flat formula Φ is *doxastically interpretable* iff it all its subformulas are inside the scope of a modality, or equivalently, if it is a Boolean combination of doxastic atoms; A *normal positive doxastically interpretable formula* (NPDI) is a formula of the form $\mathbf{K}\varphi \wedge \mathbf{B_n}\varphi_n \wedge \ldots \wedge \mathbf{B_1}\varphi_1$, where $\varphi = \varphi_\infty$, φ_1, …, φ_n are objective formulas such that for all j and $i > j$ we have $\models \varphi_j \rightarrow \varphi_i$. For instance, $\mathbf{B_4}(a \rightarrow \neg b) \vee b \wedge \mathbf{K}(a \vee c)$ is flat, but not doxastically interpretable; $\neg\mathbf{B_2}a$ and $(\mathbf{B_4}a \wedge \mathbf{B_2}\neg b) \vee \mathbf{K}(a \vee c)$ are doxastically interpretable; $\mathbf{K}(a \vee c) \wedge \mathbf{B_3}a \wedge \mathbf{B_2}a \wedge \mathbf{B_1}(a \wedge \neg c)$ is a NPDI formula; $\mathbf{K}(a \vee c)$, $\mathbf{B_2}a$ and $\mathbf{B_1}(a \wedge \neg c)$ are doxastic atoms (and *a fortiori*, NPDI formulas)[1].

We now give a simplified semantics for *flat* formulas[2]. Let $S = 2^{VAR}$ be the (finite) set of *states* (or worlds) associated with VAR.

[1] When writing NFDI formulas, we omit tautological doxastic atoms: therefore we write $\mathbf{B_2}a$ instead of $\mathbf{K}\top \wedge \mathbf{B_1}\top \wedge \mathbf{B_2}a$.

[2] This semantics can be generalized so as to interpret non-flat formulas as well; it can then be proven that any formula can be rewritten into an equivalent flat formula (as it is the case for S5 and KD45). The general semantics uses graded accessibility relations instead of OCFs and its exposition would be longer and needless for this paper.

States are denoted by s, s' etc. If φ is objective then we note $Mod(\varphi) = \{s \in S | s \models \varphi\}$. For $A \subseteq S$, $Form(A)$ is the formula (unique up to logical equivalence) such that $Mod(Form(A)) = A$.

Definition 1 (belief states) *An* ordinal conditional function (OCF) *[20], also called a* belief state, *is a function $\kappa : S \longmapsto \overline{\mathbb{N}}$ such that $\min_{s\in S} \kappa(s) = 0$. Intuitively, $\kappa(s)$ is the "exceptionality degree" of s*[3]. *In particular, $\kappa(s) = 0$ means that s is not an exceptional state while $\kappa(s) = +\infty$ means that s is impossible. The* void belief state *κ_{void} is defined by $\kappa_{void}(s) = 0$ for all s. κ is extended to objective formulas by $\kappa(\varphi) = \min\{\kappa(s) \mid s \models \varphi\}$.*

Definition 2 *A model for $\mathsf{KD45}_G$ is a pair $\langle s^*, \kappa\rangle$ where $s^* \in S$, κ an OCF, and $\kappa(s^*) < +\infty$.*

The truth of a flat formula of $\mathcal{L}$ in a state s of a model (κ, s^*) is defined by :

- for φ objective, $(\kappa, s^*) \models \varphi$ *iff* $s^* \models \varphi$;
- for φ objective and $i \in \overline{\mathbb{N}}$, $(\kappa, s^*) \models \mathbf{B_i}\varphi$ *iff* $\kappa(\neg\varphi) \geqslant i$;
- $(\kappa, s^*) \models \Phi \vee \Psi$ *iff* $(\kappa, s^*) \models \Phi$ or $(\kappa, s^*) \models \Psi$
- $(\kappa, s^*) \models \neg\Phi$ *iff* $(\kappa, s^*) \not\models \Phi$.

The connectives $\wedge, \rightarrow, \leftrightarrow$ are defined from $\vee$ and $\neg$ in the usual way. Φ is *valid* (resp. *satisfiable*) *iff* is is satisfied in any model (resp. in at least one model). Ψ is a *consequence* of Φ (denoted by $\Phi \models \Psi$) *iff* for any (κ, s^*), $(\kappa, s^*) \models \Phi$ implies $(\kappa, s^*) \models \Psi$. Φ and Ψ are equivalent *iff* $\Phi \models \Psi$ and $\Psi \models \Phi$.

In (κ, s^*), s^* represents the (objective) actual state while κ represents the agent's subjective beliefs. When Φ is doxastically interpretable, s^* has no influence on the truth value of Φ: only the subjective beliefs count (hence the terminology "doxastically interpretable"), therefore, abusing notations, for doxastically interpretable formulas we note $\kappa \models \Phi$ instead of $(\kappa, s^*) \models \Phi$.

Any belief state κ corresponds to a NPDI formula, unique up to logical equivalence:

Definition 3 (from belief states to NPDI formulas and vice versa)

1. *for any belief structure κ, $H(\kappa) = \Phi_\kappa$ is the NPDI formula (unique up to logical equivalence) defined by $\Phi_\kappa = \bigwedge \mathbf{B_i}\varphi_i$ where $\varphi_i = Form(\{s \in S \mid \kappa(s) < i\})$.*[4]
2. *given a NPDI formula $\Phi = \bigwedge_i \mathbf{B_i}\varphi_i$, $G(\Phi) = \kappa_\Phi$ is the OCF defined by*
$$\forall s \in S,\ \kappa_\Phi(s) = \min_{\kappa \models \Phi} \kappa(s)$$

For example, let κ defined by $\kappa([a, \neg b]) = 0$, $\kappa([a, b]) = 1$, $\kappa([\neg a, b]) = 1$ and $\kappa([\neg a, \neg b]) = \infty$. Then
$$H(\kappa) = \mathbf{K}(a \vee b) \wedge \mathbf{B_1}(a \wedge \neg b).$$

It is not hard to show that $G(H(\kappa)) = \kappa$ and $H(G(\Phi)) \equiv \Phi$, which means that there is a one-to-one correspondence between OCFs and equivalence classes (w.r.t. equivalence on $\mathsf{KD45}_G$) of NPDI formulas. The following properties can also be shown:

1. for any belief state κ, $H(\kappa) = \Phi_\kappa$ is the strongest (up to equivalence in $\mathsf{KD45}_G$) NPDI formula entailed by κ, i.e., for any NPDI Ψ we have $\kappa \models \Psi$ *iff* $\Phi_\kappa \models \Psi$.
2. for any NPDI formula Φ, $\kappa_\Phi = G(\Phi)$ is the minimal OCF satisfying Φ, *i.e.* ,$\kappa \models \Phi$ *iff* $\kappa \geq \kappa_\Phi$.[5]
3. for any NPDI formula $\Phi = \bigwedge \mathbf{B_i}\varphi_i$, we have $\kappa_\Phi(s) = i$ *iff* $s \models \varphi_{i+1} \wedge \neg\varphi_i$.

Notice that when writing $\Phi_\kappa = \bigwedge \mathbf{B_i}\varphi_i$, φ_i is the formula expressing *all the agent believes to the degree* i in the belief state κ. It can be shown easily that each $\mathbf{B_i}$ for $i < +\infty$ (resp. $\mathbf{K} = \mathbf{B}_\infty$) is a $\mathsf{KD45}$ (resp. $\mathsf{S5}$) modality restricted to flat formulas: thus, $\mathbf{B_i}(\varphi \wedge \psi) \leftrightarrow \mathbf{B_i}\varphi \wedge \mathbf{B_i}\psi$ and $\mathbf{K}\varphi \rightarrow \varphi$ are valid in $\mathsf{KD45}_G$. Lastly, for all i and $j \geq i$, $\mathbf{K}\varphi \rightarrow \mathbf{B_i}\varphi$ and $\mathbf{B_i}\varphi \rightarrow \mathbf{B_j}\varphi$ are valid in $\mathsf{KD45}_G$.

We now define the *combination* of belief states, and by isomorphism, the combination of NPDI formulas[6].

Definition 4 (OCF combination) *Let κ_1 and κ_2 be two OCFs. If $\min_S(\kappa_1 + \kappa_2) = \infty$, then $\kappa_1 \oplus \kappa_2$ is undefined; otherwise, $\kappa \oplus \kappa_2$ is defined by*
$$\kappa_1 \oplus \kappa_2 = \kappa_1 + \kappa_2 - \min_S(\kappa_1 + \kappa_2)$$

When defined, we have $\min_S(\kappa_1 \oplus \kappa_2) = 0$, therefore $\kappa_1 \oplus \kappa_2$ is an OCF.

By isomorphism, NPDI formulas can be combined as well:

Definition 5
$$\Phi \otimes \Psi = \begin{cases} H(\kappa_\Phi \oplus \kappa_\Psi) = H(G(\Phi) \oplus G(\Psi)) & \textit{if defined} \\ \bot & \textit{otherwise} \end{cases}$$

An important result is that if Φ, Ψ are two NPDI formulas, then $\kappa \models \Phi \otimes \Psi$ *iff* $\kappa \geqslant \kappa_\Phi \oplus \kappa_\Psi$. Moreover, the following formulas are valid:

$$\mathbf{B_i}\varphi \otimes \mathbf{B_j}\varphi \equiv \mathbf{B_{i+j}}\varphi$$
$$\mathbf{B_i}\varphi \otimes \mathbf{B_j}\neg\varphi \equiv \begin{cases} \mathbf{B_{i-j}}\varphi & \text{if } i > j \\ \mathbf{B_{j-i}}\neg\varphi & \text{if } j > i \\ \top & \text{if } i = j \end{cases};$$
$$\Phi \otimes \Psi \equiv \Psi \otimes \Phi; \quad \Phi \otimes (\Psi \oplus \Xi) \equiv (\Phi \otimes \Psi) \otimes \Xi);$$
$$\Phi \otimes \top \equiv \Phi; \quad \Phi \otimes \bot \equiv \bot$$

Importantly, $\Phi \otimes \Phi$ is generally not equivalent to Φ.

3 Observations and revision

We now consider a finite set ACT of actions available to the agent. Although an action may generally both have "physical" (or *ontic*) effects on the state of the world and give some feedback leading the agent to revise her belief state, we assume without loss of generality that ACT is partitioned into two disjoint subsets ACT_P (pure physical actions, feedback-free), and ACT_E (pure sensing actions, with no physical effects).

The feedback of a sensing action is an *observation*. Ideally, an observation is precise and reliable, but this is far from being always the case in practice. We first have to make clear what we mean by observation.

Definition 6 *An* observational belief state (OBS), *or, for short, an* observation, *is a belief state κ_{obs}, corresponding to a NPDI-formula $\Phi_{obs} = H(\kappa_{obs}) = \mathbf{K}o \wedge \mathbf{B_n}o_n \wedge \ldots \wedge \mathbf{B_1}o_1$ (by convention we write $o_\infty = o$).*

[3] $\kappa(s)$ is usually interpreted in terms of infinitesimal probabilities: $\kappa(s) = k < +\infty$ means $prob(s) = o(\varepsilon^k)$, where ε is infinitely small

[4] Recall that $\mathbf{K} = \mathbf{B}_\infty$.

[5] This property is nothing but a rewriting of the minimum specificity principle in possibilistic logic. See [9] for a discussion on the translations between OCFs and possibility distributions.

[6] Note that is not a real connective, because it only connects NPDI formulas.

An observation is therefore defined by the belief state it conveys (which, in practice, may be a function of the belief state of the source and the belief that the agent has on the reliability of the source). κ_{obs} can also be viewed as the belief state the agent is if she gets this observation in the void belief state κ_{void}.

If $\Phi_{obs} \equiv \mathbf{K}o$ then obs is a *reliable observation*, which brings a totally certain information about the actual state of the world. More generally, a simple observation $\Phi_{obs} \equiv \mathbf{B_k}o_k$, induces a belief in a single fact o_k, with a reliability degree k. In practice, many observations will be simple, but not all.

This rather complex definition is due to the fact that a single observation generally relates to the real state of the world in several ways, with various degrees of uncertainty (exactly as in the Bayesian case). Consider for instance reading the value θ on a temperature sensor, which may for instance correspond to the observation $\Phi_{obs} = \mathbf{B_1}(t-1 \leq \theta \leq t+1) \wedge \mathbf{B_2}(t-2 \leq \theta \leq t+2) \wedge \mathbf{K}(t-5 \leq \theta \leq t+5)$.

We now define how the agent revises her current belief state after an observation.

Definition 7 *Let κ be a belief state and κ_{obs} an observational belief state. The revision of κ by κ_{obs} is simply the combination of κ and κ_{obs},* i.e. , $rev(\kappa, o) = \kappa \oplus \kappa_{obs}$.

In particular, when $obs = \mathbf{K}o$ is a reliable observation, we get

$$(\kappa \oplus \kappa_{obs})(s) = \kappa(s|obs) = \begin{cases} +\infty & \text{if } s \models \neg o \\ \kappa(s) - \kappa(o) & \text{if } s \models o \end{cases}$$

that is, a conditioning by o in the sense of [20]. By isomorphism, revision can be performed syntactically: $\Phi = \mathbf{K}\varphi \wedge \mathbf{B_n}\varphi_n \wedge \ldots \wedge \mathbf{B_1}\varphi_1$ being a NPDI formula and $obs = \mathbf{K}o \wedge \mathbf{B_n}o_n \wedge \ldots \wedge \mathbf{B_1}o_1$ an observation, the revision Φ by obs is $\Phi \otimes obs$. The following result shows how the latter expression can be computed syntactically in a compact way, *without performing revision state by state*:

Proposition 1

$$\Phi \otimes \Phi_{obs} \equiv \mathbf{B_1}\psi_p \wedge \ldots \wedge \mathbf{B_m}\psi_{p+m-1} \wedge \mathbf{K}\psi$$

where

- $\psi = \varphi \wedge o$
- $\psi_i = (\varphi_1 \wedge o_i) \vee (\varphi_2 \wedge o_{i-1}) \vee \ldots \vee (\varphi_i \wedge o_1)$;
- $p = \min\{j, \psi_j \not\equiv \bot\}$;
- $m = \min\{k, \psi_{p+k} \equiv \psi\}$.

Example 1 *The agent wants to know the direction to the railway station. Assume there are only two directions, r (right) and $\neg r$ (left). His initial belief state is void ($\kappa_0 = \kappa_{void}$). He tries to acquire some more information by asking pedestrians. Let us consider five possible observations: $\Phi_{obs_1} = \mathbf{B_2}r$ ("the station is on the right", given without hesitation, $\Phi_{obs_2} = \mathbf{B_1}r$ ("I believe it's on the right but I'm not sure"), $\Phi_{obs_3} = \mathbf{B_2}\neg r$, $\Phi_{obs_4} = \mathbf{B_1}\neg r$ and finally $\Phi_{obs_5} = \mathbf{K}\top$ ("I have no idea"). After observing first obs_2, we have $\kappa_1 = \kappa_0 \oplus \kappa_{obs_2} = \kappa_{obs_2}$ and $\Phi_1 = \Phi_0 \oplus \Phi_{obs_2} = \mathbf{B_1}r$. Assume now that the second pedestrian gives obs_2 too. Using Proposition1, we get: $\psi = \top \wedge \top = \top$; $\psi_1 = r \wedge r = r$; $\psi_2 = (r \wedge \top) \vee (\top \wedge r) = r$; $\psi_3 = (r \wedge \top) \vee (\top \wedge \top) \vee (\top \wedge r) = \top$ – therefore $p = 1$ and $p + m = 3$, hence $\Phi_1 \otimes \Phi_{obs_1} = \mathbf{B_2}r \wedge \mathbf{K}\top = \mathbf{B_2}r$. If the second observation had been obs_4 instead of obs_2 we would have had $\Phi_2 = \Phi_1 \otimes \Phi_{obs_4} = \mathbf{K}\top$ (the agent comes back to his initial belief state). Indeed, $\psi = \top \wedge \top = \top$; $\psi_1 = r \wedge \neg r = \bot$; $\psi_2 = (r \wedge \top) \vee (\top \wedge \neg r) = \top$ – therefore $p = 2$ and $p + m = 2$, hence $\Phi_1 \otimes \Phi_{obs_2} = \mathbf{B_1}\top \wedge \mathbf{K}\top = \mathbf{K}\top$. By induction, one can show that after p_1 occurrences of obs_1, p_2 of obs_2, p_3 of obs_3, p_4 of obs_4 and p_5 of obs_5 (in any order), iterated combination leads to*

- $\mathbf{B_q}r$ *if* $2p_1 + p_2 > 2p_3 + p_4$ *and* $q = (2p_1 + p_2) - (2p_3 + p_4)$;
- $\mathbf{B_q}\neg r$ *if* $2p_1 + p_2 < 2p_3 + p_4$ *and* $q = (2p_3 + p_4) - (2p_1 + p_2)$;
- $\mathbf{K}\top$ *if* $2p_1 + p_2 = 2p_3 + p_4$.

This example shows how observations *reinforce* prior beliefs when they are consistent with them[7]. It clearly appears that the crucial hypothesis underlying the combination rule is *independence* between the successive observations. Thus, on Example 1, the successive answers are independent (pedestrians do not listen to the answers given by their predecessors)[8].

4 Progression

Physical actions may change the state of the world but do not give any feedback. Therefore, given an initial belief state κ and an ontic action α, it is possible to determine the future belief state (after the action is performed) by projecting the possible outcomes of α on the current belief state. This operation is usually called *progression*: $prog(\kappa, \alpha)$ is the belief state obtained after α is performed in belief state κ. By isomorphism, we also define $Prog(\Phi, \alpha)$ where $Prog(\Phi, \alpha) = H(prog(G(\Phi), \alpha))$.

The semantics of progression is defined as in [5] by means of *OCF transition models*.

Definition 8 *An OCF transition model is a collection of OCFs $\{\kappa_\alpha(.|s), s \in S\}$.*

$\kappa_\alpha(s'|s)$ is the exceptionality degree of the outcome s' when performing action α in state s. Notice that for all $s \in S$, $\kappa_\alpha(.|s)$ is an OCF, which means that $\min_{s' \in S} \kappa_\alpha(s'|s) = 0$.

Definition 9 (progression of κ by an ontic action) *Given an initial belief state κ and an ontic action α whose dynamics is expressed by the OCF transition model κ_α, the progression of κ by α is the belief state $\kappa' = prog(\kappa, \alpha)$ defined by*

$$\forall s' \in S,\ \kappa'(s') = \min_{s \in S} \{\kappa(s) + \kappa_\alpha(s'|s)\}$$

Notice that κ' is a belief state, because the normalization of *both* κ and $\kappa_\alpha(.|s)$ implies that $\min_{s' \in S} \{\min_{s \in S} \{\kappa(s) + \kappa_\alpha(s'|s)\}\} = 0$, *i.e.* $min\{\kappa'(s'),\ s' \in S\} = 0$.

Example 2 *Consider two blocks A and B initially put down on a table; the propositional variable x is true if A is on top of B, false otherwise. A robot can perform the action α consisting in try to put A on B. If A is on B in the initial state, the action has no effect; otherwise, it normally succeeds (ie, x becomes true), and exceptionally fails (in that case, x remains false). The OCF transition model for α is: $\kappa_\alpha(x|x) = 0$; $\kappa_\alpha(\neg x|x) = \infty$; $\kappa_\alpha(x|\neg x) = 0$; $\kappa_\alpha(\neg x|\neg x) = 1$. If the initial state is κ_{void}, then $prog(\kappa_{void}, \alpha) = \{(x, 0), (\neg x, 1)\}$ and $prog(\kappa', \alpha) = \kappa'' = \{(x, 0), (\neg x, 2)\}$. More generally, after*

[7] This has to be contrasted with *transmutations* [23], where one enforces the new belief state to satisfy a constraint of the form $\kappa(\varphi) = i$. Probability theory has also both kinds of rules: Jeffrey's (without implicit reinforcement) and Pearl's (see [7] for a discussion).

[8] If this assumption could no longer be made, then we should have to model it by adding a new variable which would have the effect of blocking (or limiting) the reinforcement.

performing α n times without performing any sensing action (starting from κ_{void}), we get $prog(\kappa_{void}, \alpha^n) = \{(x, 0), (\neg x, n)\}$, whose associated NPDI formula is $\mathbf{B_n}x$: the agent believes to the degree n that A is on B.

Example 2 shows that once again, the underlying hypothesis is the independence between the outcomes of the different occurrences of actions. Indeed, the intuitive explanation of the result of previous example is that after these n executions of α, A is still not on B if and only if all n occurrences of α failed; each of the failures has an excaptionality degree of 1 and failures are independent, henceforth, n successive failures occur with an exceptionality degree of n. Notice that this reinforcement effect is a consequence of the use of $\oplus$ (if conjunction were used instead, we would still get $\mathbf{B_1}x$ after performing α n times).

We now show how progression can be computed syntactically, which avoids explicitly computing progression state by state consisting of a straightforward application of the definition.

First, we assume that the effects of actions are described by a *graded action theories*, generalizing action theories so as to allow for more or less exceptional action effects. In order to do so, we extend (as usual) the initial propositional language $\mathcal{L}$ by duplicating each variable x of VAR in x_t and x_{t+1} (representing x respectively before and after the execution of the action); let $VAR_t = \{x_t | x \in VAR\}$ and $VAR_{t+1} = \{x_{t+1} | x \in VAR\}$, $S_t = 2^{VAR_t}$ and $S_{t+1} = 2^{VAR_{t+1}}$. For any formula Φ, Φ_t (resp. Φ_{t+1}) denotes the formula obtained by replacing each occurrence of x by x_t (resp. x_{t+1}). A *gradual action theory* is a NPDI formula of this extended language: $\Sigma_\alpha = \mathbf{K}r \wedge \mathbf{B_n}r_n \wedge \ldots \wedge \mathbf{B_1}r_1$. The gradual action theory is obtained from a set of causal (dynamic or static) rules through a completion process whose technical details are omitted due to lack of space. Notice however that this completion does not present any particular difficulty: it is an easy extension of completion for nondeterministic action theories such as in [16, 11]. We just give the graded action theory corresponding to Example 2:

$$\Sigma_\alpha = \mathbf{K}(x_t \to x_{t+1}) \wedge \mathbf{B_1}x_{t+1}$$

Like for the static case, any OCF transition models correspond to graded action theories and *vice versa*: $\{\kappa_\alpha(.|s), s \in S\}$ induces $\Sigma_\alpha = \mathbf{K}r \wedge \mathbf{B_n}r_n \wedge \ldots \wedge \mathbf{B_1}r_1$ where $r_i = Form\{(s'_{t+1}, s_t) \mid \kappa_\alpha(s'_{t+1}|s_t) < i\}$.

Now, we recall the definition of *forgetting* a subset of propositional variables X from an objective propositional formula ψ [17]:

1. $forget(\{x\}, \psi) = \psi_{x \leftarrow \top} \vee \psi_{x \leftarrow \bot}$;
2. $forget(X \cup \{x\}, \psi) = forget(\{x\}, forget(X, \psi))$.

Forgetting is extended to S5 formulas in [15], and is here extended to NPDI formulas in the following way: if $\Phi = \mathbf{K}\varphi \wedge \mathbf{B_n}\varphi_n \ldots \wedge \mathbf{B_1}\varphi_1$ and $X \subset Var(\Phi)$, then $Forget(X, \Phi) = \mathbf{K}forget(X, \varphi) \wedge \mathbf{B_n}forget(X, \varphi_n) \wedge \ldots \wedge \mathbf{B_1}forget(X, \varphi_1)$.

Now we have the following syntactical characterization of progression:

Proposition 2 *Let Φ be the NPDI formula corresponding to the initial belief state κ, and α an ontic action described by an action theory as previously defined. Then*

$$Prog(\Phi, \alpha) \equiv Forget(VAR_t, \Phi_t \otimes \Sigma_\alpha)$$

Thus, progression amounts to a combination followed by a forgetting. For the first step, Proposition 1 can be applied again, as shown on the following example. The second step amounts to a sequence of classical forgetting operations.

Example 2 (continued) *We have $\Sigma_\alpha = \mathbf{K}(x_t \to x_{t+1}) \wedge \mathbf{B_1}x_{t+1}$. The initial belief state corresponds to $\Phi = \mathbf{B_1}x$. Then, $\Phi_t \otimes \Sigma_\alpha = \mathbf{K}\psi \wedge \mathbf{B_n}\psi_n \wedge \ldots \wedge \mathbf{B_1}\psi_1$, where $\psi = \top \wedge (x_t \to x_{t+1})$; $\psi_1 = x_t \wedge x_{t+1}$; $\psi_2 = (x_t \wedge (x_t \to x_{t+1})) \vee (\top \wedge x_{t+1})$; $\psi_3 = (x_t \wedge (x_t \to x_{t+1})) \vee (\top \wedge (x_t \to x_{t+1})) \vee (\top \wedge x_{t+1})$. After simplifying the expression we get $\psi = x_t \to x_{t+1}$; $\psi_1 = x_t \wedge x_{t+1}$; $\psi_2 = x_{t+1}$; $\psi_3 = x_t \to x_{t+1} = \psi$. Next we get $\Phi_t \otimes \Sigma_\alpha \equiv \mathbf{K}(x_t \to x_{t+1}) \wedge \mathbf{B_1}(x_t \wedge x_{t+1}) \wedge \mathbf{B_2}x_{t+1}$ and $Forget(VAR_t, \Phi_t \otimes \Sigma_\alpha) = \mathbf{K}\top \wedge \mathbf{B_1}x_{t+1} \wedge \mathbf{B_2}x_{t+1} \equiv \mathbf{B_2}x_{t+1}$, and finally $Prog(\Phi, \alpha) = \mathbf{B_2}x$.*

Note the importance of combination, which explains the reinforcement obtained[9].

5 Belief-based programs

A belief-based program is built up from the set of primitive actions ACT and usual program constructor. Given a set $ACT = ACT_P \cup ACT_E$ of primitive actions, a *belief-based program* (BBP) is defined inductively as follows:

- the empty plan λ is a BBP;
- for any $\alpha \in ACT$, α is a BBP;
- if π and π' are BBP then $\pi; \pi'$ is a BBP;
- if π and π' are BBP and Φ is a NPDI formula, then `if` Φ `then` π `else` π' and `while` Φ `do` π are BBPs.

Thus, a BBP is a program *whose branching conditions are doxastically interpretable*: the agent can decide whether she *believes* to a given degree that a formula is true (whereas she is generally unable to decide whether a given objective formula is true in the actual world). For instance, the agent performing the BPP

$$\begin{array}{rl} \pi = & \texttt{while } \neg(\mathbf{B_2}r \vee \mathbf{B_2}\neg r) \texttt{ do } ask; \\ & \texttt{if } \mathbf{B_2}r \texttt{ then } goright \texttt{ else } goleft \end{array}$$

performs the sensing action ask until she has a belief firm enough (namely of degree 2) about the way to follow (notice that this program is not guaranteed to stop!).

The execution of a belief program is a *nondeterministic* function mapping a pair consisting of an initial belief state and a program to a sequence of actions and observations. Each sensing action α is associated with a *nondeterministic* function $feedback : ACT_E \to OBS$. $feedback(\alpha)$ is the observation obtained after performing the sensing action α in a given environment. If α is a precise and reliable truth test (that is, α senses the value of a fluent f), then $feedback(\alpha) \in \{obs(f), obs(\neg f)\}$ – where $obs(f) = \mathbf{K}f$ and $obs(\neg f) = \mathbf{K}\neg f$). But generally, there may be any number of possible outcomes for a given sensing action, including possible void observations ($obs(void) = \mathbf{K}\top$). Since off-line reasoning is outside the scope of this paper, it is needless to specify formally which observations are possible (and how likely they are in a given belief state) after a given sensing action is performed.

[9] Such a reinforcement would not be obtained if conjunction were used instead of combination: doing α many times would give $\mathbf{B_1}x$ again and again.

The execution of a BPP π in an initial belief state κ is defined as a nondeterministic notion (that is, there is generally a set of possible executions of π in κ).

Definition 10 *Given a BPP π and a belief state κ, the execution of π in κ is the nondeterministic function defined inductively by*

- $exec(\lambda, \kappa) = stop$;
- *if* $\pi = \alpha; \pi'$ *with* $\alpha \in ACT_P$ *then* $exec(\pi, \kappa) = do(\alpha); exec(\pi', prog(\kappa, \alpha))$
- *if* $\pi = \alpha; \pi'$ *with* $\alpha \in ACT_E$ *then* $exec(\pi, \kappa)$ $= do(\alpha); obs := feedback(\alpha); exec(\pi', rev(\kappa, obs))$
- *if* $\pi = \texttt{if } \Phi \texttt{ then } \pi' \texttt{ else } \pi''$ *then* $exec(\pi, \kappa) = \begin{cases} exec(\pi', \kappa) & \text{if } \kappa \models \Phi \\ exec(\pi'', \kappa) & \text{otherwise} \end{cases}$
- *if* $\pi = (\texttt{while } \phi \texttt{ do } \pi'); \pi''$ *then* $exec(\pi, \kappa) = \begin{cases} exec((\pi'; \pi), \kappa) & \text{if } \kappa \models \Phi \\ exec(\pi'', \kappa) & \text{otherwise} \end{cases}$

Therefore, the execution of a program in a belief state is a *nondeterministic* function whose output (or trace) is a sequence of action execution instructions $do(\alpha)$ and observations outcomes $obs := feedback(\alpha)$. For instance, a possible execution of the program π above in the initial belief state κ_{void} is $\langle ask; obs(\mathbf{B_1}\neg r); ask; obs(\mathbf{B_1}\neg r); do(goleft)\rangle$.

A further formalization, along the lines followed by the formal modeling of on-line execution of Golog programs [8, 13], is a topic for further study.

6 Related work

6.1 Graded belief

The construction given in Part 2 is not highly original (and is not the primary goal of the paper). It is very similar to the work on stratified belief bases and possibilistic logic where the duality between (semantical) belief states and (syntactical) NPDI formulas, as well as combination operators, exists under a similar form [4]. As to gradual doxastic logics, [22] who define a gradual version of KD45 somewhat differet from ours, where $\mathbf{B_n}\varphi$ expresses that φ is true in all worlds except n or less.

6.2 Revision by uncertain observations

The closest work to ours is [6], where observational systems allowing for unreliable observations are modelled using OCFs. Their work is less specific than ours (notice that in the absence of ontic actions, our revision process falls in the the class of Markovian observation systems). The main difference between [6] and our Section 3 is that the revision functions in [6] remain defined at the semantical level, which, if computed state by state following the definition, needs an exponantially large data structure. Our approach can therefore be viewed as providing a compact representation for a specific class of observation systems. In another line of work, namely [1], models noisy observations in a probabilistic version of the situation calculus (again, compact representation issues are not considered). Belief transmutations and adjustments [23] are based on OCFs too; however, like Jeffrey's rule in probability theory, they consist in changing minimally a belief state so as to force a given formula to have a given exceptionality degree, which therefore drastically differs from a revision rule enabling an *implicit* reinforcement of belief when the observation is consistent with the initial belief state, as seen in Example 1. See also [9] for a panorama of revision rules in numerical formalisms, including OCFs.

6.3 Actions with exceptional effects

[12] and [5] study belief update operators with belief states modelled by OCFs, so as to model exceptional effects of actions. These operators are very similar to our progression for ontic actions from a semantical point of view – but they do not give any syntactical characterization of progression. [19] considers physical and sensing actions in a situation calculus setting, where states are mapped to a plausibility values; these plausibility values are simply inherited from plausibility values in the initial belief state (noisy observations and exceptional effects actions are not considered).

REFERENCES

[1] F. Bacchus, J. Halpern, and H. Levesque. Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence*, 111:171–208, 1999.

[2] F. Bacchus and R. Petrick. Modelling an agent's incomplete knowledge during planning and execution. In *KR-98*, pages 432–443, 1998.

[3] C. Baral and T. Son. Formalizing sensing actions – a transition function based approach. *Artificial Intelligence*, 125(1–2):19–91, 2001.

[4] S. Benferhat, D. Dubois, and H. Prade. A computational model for belief change and fusing ordered belief bases. In *Frontiers in Belief Revision (M.A. Williams and H. Rott, eds.)*, pages 109–134. Kluwer Academic Publishers, 2001.

[5] C. Boutilier. A unified model of qualitative belief change: A dynamical systems perspective. *Artificial Intelligence Journal*, 98(1-2):281–316, 1998.

[6] C. Boutilier, N. Friedman, and J. Halpern. Belief revision with unreliable observations. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)*, pages 127–134, 1998.

[7] H. Chan and Darwiche A. On the revision of probabilistic beliefs using uncertain evidence. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI-03)*, 2003.

[8] G. De Giacomo and H. Levesque. An incremental interpreter for high-level programs with sensing. In *Logical Foundation for Cognitive Agents: Contributions in Honor of Ray Reiter*, pages 86–102. Springer, Berlin, 1999.

[9] D. Dubois and H. Prade. A synthetic view of belief revision with uncertain inputs in the framework of possibility theory. In *Int. J. of Approximate Reasoning*, volume 17, pages 295–324. 1997.

[10] R. Fagin, J. Halpen, Y. Moses, and M. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.

[11] E. Giunchiglia, J. Lee, N. McCain, V. Lifschitz, and H. Turner. Nonmonotonic causal theories. *Artificial Intelligence*, 153:49–104, 2003.

[12] M. Goldszmidt and J. Pearl. Rank-based systems: A simple approach to belief revision, belief update, and reasoning about evidence and actions. In *Proceedings of KR'92*, pages 661–672, 1992.

[13] H. Grosskreutz and G. Lakemeyer. On-line execution of ccgolog plans. In *Proc. IJCAI-01*, pages 12–18, 2001.

[14] A. Herzig, J. Lang, D. Longin, and Th. Polacsek. A logic for planning under partial observability. In *AAAI-00*, pages 768–773, 2000.

[15] A. Herzig, J. Lang, and P. Marquis. Action representation and partially observable planning in epistemic logic. In *Proceedings of IJCAI03*, pages 1067–1072, 2003.

[16] F. Lin. Embracing causality in specifying the indeterminate effects of actions. In *Proc. of AAAI'96*, 1996.

[17] F. Lin and R. Reiter. Forget it! In *Proceedings of the AAAI Fall Symposium on Relevance*, pages 154–159, New Orleans, 1994.

[18] R. Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT Press, 2001.

[19] S. Shapiro, M. Pagnucco, Y. Lesperance, and H. Levesque. Iterated belief change in the situation calculus. In *Proceedings ok KR2000*, pages 527–537, 2000.

[20] W. Spohn. Ordinal conditional functions: a dynamic theory of epistemic states. In William L. Harper and Brian Skyrms, editors, *Causation in Decision, Belief Change and Statistics*, volume 2, pages 105–134. Kluwer Academic Pub., 1988.

[21] M. Thielscher. Programming of reasoning and planning agents with FLUX. In *Proceedings of KR2002*, pages 435–446. Morgan Kaufmann, 2002.

[22] W. van der Hoek and J.-J.Ch. Meyer. Graded modalities for epistemic logic. *Logique et Analyse*, 133-134:251–270, 1991.
[23] M.-A. Williams. Transmutations of knowledge systems. In *Proceedings of KR'94*, pages 619–629, 1994.