

Dynamic Epistemic Logics

Jan van Eijck

Abstract Dynamic epistemic logic, broadly conceived, is the study of rational social interaction in context, the study, that is, of how agents update their knowledge and change their beliefs on the basis of pieces of information they exchange in various ways. The information that gets exchanged can be about what is the case in the world, about what changes in the world, and about what agents know or believe about the world and about what others know or believe. This chapter gives an overview of dynamic epistemic logics, and traces some connections with propositional dynamic logic, with planning and with probabilistic updating.

1 Introduction

Logic is broadly conceived in Johan van Benthem's work as the science of information processing in evolving contexts. In most logic textbooks, with [18] as a notable exception, it is assumed that the reasoning processes that constitute the subject matter of the discipline take place in the head of an ideal reasoner. The validity of an argument, such as a step of Modus Ponens establishes a fact with the help of another fact plus an implication, and the agent performing these steps is kept out of the picture.

But in fact, the pieces of information that are put together by means of applications of Modus Ponens can have many different sources, involving different agents, and communication between them. Here is an early Chinese example that Johan is fond of quoting:

Someone is standing next to a room and sees a white object outside. Now another person tells her that there is an object inside the room of the same colour as the one outside. After all this, the first person knows that there is a white object inside the room. This is based on

Jan van Eijck
CWI, Science Park 123, 1098 XG Amsterdam and ILLC, Science Park 107, 1098 XG Amsterdam,
e-mail: jve@cwi.nl

three actions: an observation, then an act of communication, and finally an inference putting things together.

To give a full account of what goes on here, one has to represent the state of knowledge (or belief) of several agents, and model what goes on when they perceive facts in the world, when information is exchanged between them, and when they act on the world by make changes to it. This is what dynamic epistemic/doxastic logic is all about.

2 Knowledge, Belief and Change

The original account of belief and knowledge in terms of possible states of affairs or possible worlds is due to Hintikka [43], who proposed to analyze knowledge and belief with the tools of modal logic. Knowing about a fact p boils down to the ability to distinguish states of affairs where p is true from states of affairs where p is false, and the key notion of epistemic logic is that of an *indistinguishability relation* between possible worlds.

This analysis was taken up by cognitive scientists [34], computer scientists [39, 31, 30] and game theorists [4, 12, 58], and gradually extended to include interaction between different agents. It turned out that the notion of *common knowledge* plays a key part in the analysis of rational behaviour in games.

Dynamic epistemic/doxastic logic (see [22] for a textbook treatment) studies the evolution of knowledge and belief in the context of change. This change can be of various kinds:

- Changing beliefs about an unchanging world: in the Chinese room example the world does not change, but the first person learns something new about what is the case.
- Changing beliefs about a changing world: imagine a robot finding its way through a maze. The robot moves through the maze to a different spot and observes what it finds there. The change of location is a change in the states of affairs in the real world, the new observation causes the robot to change its belief state.
- Incorporating or failing to incorporate information about change in the world: a voter is taking part in an election process, but misses the communication about a change in the rules of the voting game.

In epistemic logic this is all expressed qualitatively, but there is an obvious relation to numerical ways of expressing rational belief. A change in belief could also be a change in the probability estimation that something is the case. Changes in the world may be thought of as being the result of indeterminate actions that occur with a certain probability. See below, Section ??, for probabilistic extensions of DEL that can cover such cases.

The distinction between qualitative ways and quantitative ways of expressing preference can also be found in game theory, where abstract strategic games are

expressed in terms of individual preference relations on outcomes, while concrete strategic games use payoff functions that represent these preferences [55], and in probability theory, where the well-known Cox axioms list three conditions on ‘degrees of belief’ that are sufficient for a map to quantitative probabilities [21].

3 The Dynamic Turn in Epistemic/Doxastic Logic

A pioneer paper shifting the focus from the study of information states *simpliciter* to information change is the work of Jan Plaza on the connection between public announcement and the generation of common knowledge [59]. This was followed up in a number of PhD theses under the supervision of Johan van Benthem: [45], [38], [35].

In an ILLC report from 2000, Johan van Benthem analyzes the kind of information update by means of world elimination that goes on in public announcement as model relativization [17], connecting truth in a full model with truth in a model relativized by some unary predicate. This explains why epistemic logic with public announcement can be axiomatized by means of reduction axioms that spell out the recursive definition of the relativization operation. The public announcement logic of epistemic logic with a common knowledge operator added admits no such axiomatization: the recursive definition of restriction on a common knowledge formula proceeds via relativized common knowledge.

A next key contribution is the proposal to view information updates themselves as Kripke models representing the agent’s take on what happens when information is updated [7, 8, 6]. This generalizes information updating to the whole class of multiagent Kripke models, not just multi-agent S5 models (an S5 model is a Kripke model where all accessibility relations are equivalence relations).

Action models are in fact epistemic/doxastic perspectives on communicative events. This epistemic/doxastic perspective on communication had already emerged in the AI literature, where epistemic reasoning was integrated into the situation calculus (essentially, a version of first order logic designed to describe changes in states of affairs) by Bob Moore [54]. The study of noisy sensors in [5] provides an epistemic perspective on communication in the situation calculus, by analyzing noisy observations as epistemic update events (S5 update models, in fact).

While action model update works for all Kripke models, multi-agent S5 updating of multi-agent S5 models is an important special case. The product of an S5 model with an S5 action model is again S5. The reason for this is that the S5 properties of reflexivity, transitivity and euclidity are preserved under the update operation of restriction combined with product.

A sentence of first order logic is preserved under restriction and product if and only if the sentence is universal Horn, where a universal Horn sentence of FOL is the universal closure of a disjunction with at most one atom disjunct, and with the remaining disjuncts negations of atoms [52]. Reflexivity, transitivity and euclidity are universal Horn, but seriality is not. This explains why the update of a KD45 epis-

temic model with a KD45 action model (a model where the accessibility relations are transitive, euclidi and serial) may yield a result that is not KD45. Information update of belief models with belief action models has a glitch.

This was one of the motivations to explore combinations of information update with information change as belief upgrade, in [2], followed by [15] and [9, 10, 11]. The information update format was extended still further in [37].

Taking an extensional view on the update mechanisms involved in knowledge update, belief revision, belief change, and factual change, it turns out that these can all be described in terms of PDL style operations. Indeed, all of these update formats can be captured in a general logic of communication and change, and [16] proves a technical result that adding such update mechanisms to the logic of epistemic PDL (where the basic actions describe primitive epistemic/doxastic relations, and the PDL operators construct more complex epistemic/doxastic relations) does not increase the expressive power of the logic. This inspired further work on how PDL can be viewed (or: reinterpreted) as a logic of belief revision, in [28], and as a multi-agent strategy logic, in [27].

However, this cannot be the whole story. This extensional view may illuminate the bare bones of DEL, but it disregards the flesh. Updating is a process where agents may follow specific update protocols, and it makes eminent sense to study possible formats of update rules. Such connections with protocol dynamics were explored in [64], and they led to interesting work on the use of DEL in planning [1].

Just as in matters of computation, it pays off to shift from an extensional to an intensional view. Extensionally, all we can say about the objects of computation is that they are the recursive functions. Intensionally, we can say a lot more, by focussing on the *how* of computation. Similarly with communication: the extensional view disregards the inner workings of *how* information gets processed, and an intensional view on DEL brings this to light.

4 Announcements and Updating

The cartoon in Figure 1 illustrates what goes on when public announcements are processed by perfectly rational agents, in this case logicians looking for something to drink. It is assumed that they all know what they want to drink themselves, but are uncertain about the wishes of the others. The cartoon tells the story of what happens then.

The question “Does everyone want beer?” triggers the following instruction, say for agent i (we use b_i for “ i wants beer” and \Box_i for “ i knows”, plus the usual boolean connectives):

- If $\Box_i(b_1 \wedge b_2 \wedge b_3)$ then i says “Yes”.
- If $\Box_i\neg(b_1 \wedge b_2 \wedge b_3)$ then i says “No”.
- Otherwise, i says “I don’t know”.

These answers themselves serve as updates:




Fig. 1 Three logicians, from <http://spikedmath.com>

- i says “Yes”: update with public announcement of $\Box_i(b_1 \wedge b_2 \wedge b_3)$
- i says “No”: update with public announcement of $\Box_i\neg(b_1 \wedge b_2 \wedge b_3)$.
- i says “I don’t know”: update with public announcement of

$$\neg\Box_i(b_1 \wedge b_2 \wedge b_3) \wedge \neg\Box_i\neg(b_1 \wedge b_2 \wedge b_3).$$

The updates are instructions to *eliminate* worlds. The update with $\neg\Box_i(b_1 \wedge b_2 \wedge b_3) \wedge \neg\Box_i\neg(b_1 \wedge b_2 \wedge b_3)$ eliminates all worlds where $\Box_i(b_1 \wedge b_2 \wedge b_3)$ or $\Box_i\neg(b_1 \wedge b_2 \wedge b_3)$ holds.

To check a $\Box_i\phi$ formula, one has to check whether ϕ holds in all i -accessible worlds. If 1, 2, 3 are the three logicians standing or sitting next to each other in left-to-right order, then a situation where 1 wants beer and 2, 3 do not can be represented as $\bullet\circ\circ$. The space of possibilities is given by:

$$\{\circ\circ\circ, \circ\circ\bullet, \circ\bullet\circ, \bullet\circ\circ, \bullet\bullet\circ, \bullet\circ\bullet, \circ\bullet\bullet, \bullet\bullet\bullet\}.$$

How about the epistemic accessibilities? These accessibilities should express that each agent initially only knows about her own state (thirsty for beer or not). If — represents the accessibility of the first logician, then the state $\circ\circ\circ$ should be connected by a — line to $\circ\circ\bullet$, to $\circ\bullet\circ$, and to $\bullet\circ\circ$ (all states where the other logicians have different wants). Thus, initially, the space of possibilities together with the epistemic accessibilities is given by the picture in Figure 2, with \bullet for “wants beer”, solid lines for 1, dashed lines for 2 and dotted lines for 3, where 1,2,3

are the three logicians standing or sitting next to each other. Note that states s and t in the picture are linked by lines for agent i iff either in both of s, t agent i wants beer or in both of s, t agent i does not want beer.




Fig. 2 Initial situation when three logicians enter a bar

After the first logician says “I don’t know”, the possibilities where

$$\Box_i(b_1 \wedge b_2 \wedge b_3) \text{ or } \Box_i\neg(b_1 \wedge b_2 \wedge b_3)$$

are true drop out. Note that these are precisely the possibilities where she does not want a beer herself (Figure 3).

After the second logician says “I don’t know”, all remaining possibilities where she does not want beer drop out: see Figure 4. Now the third logician resolves the case by saying either “yes” or “no”, and if the bartender is also a perfect logician, she knows in each case which of her customers to serve a beer.

The update process in the familiar muddy children example, where perfectly rational children deduce from the assumption that at least one child has a muddy forehead, and from the mud they see or fail to see on the foreheads of the other children whether they themselves are muddy, is similar. Only the accessibility relations are the converse of those in the thirsty logicians example: the thirsty logicians know what they want to drink but do not know what the others want, while the




Fig. 3 After the announcement of the first logician




Fig. 4 After the announcement of the second logician

muddy children know about the muddiness of the others but not about their own muddiness.

5 Kripke Models and Action Model Update

The information update processes in the case of the three thirsty logicians, or of the n muddy children, are special cases of a general procedure for updating epistemic models with action models due to Baltag, Moss, Solecki [7]. This handles the communication between the drinking logicians, the muddy children, and much more besides.

Let a finite set Ag of agents and a set $Prop$ of propositions be given. Then the class of Kripke models over Ag and $Prop$ is given by:

Definition 1. A **Kripke Model** is a tuple (W, R, V) where

- W is a non-empty set of worlds.

- R is a function that assigns to every agent $a \in A$ a binary relation R_a on W .
- V is a valuation function that assigns to every $w \in W$ a subset of $Prop$.

An action model is like a Kripke model for Ag and $Prop$, with the difference that the worlds are now called *actions* or *events*, and that the valuation has been replaced by a map **pre** that assigns to each event e a formula of a suitable epistemic language called the *precondition* of e . Let us fix the language first:

Definition 2. The multimodal language \mathcal{L} over Ag and $Prop$ is given by the following BNF definition, where a ranges over Ag and p over $Prop$:

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \phi \mid \Box_a \phi.$$

We assume the usual abbreviations for \vee , \rightarrow , \leftrightarrow , \Diamond_a .

The truth definition for this language is given by:

Definition 3. Let $\mathbf{M} = (W, R, V)$ and $w \in W$ in:

$$\begin{aligned} \mathbf{M} \models_w p &\text{ iff } p \in V(w) \\ \mathbf{M} \models_w \neg\phi &\text{ iff it is not the case that } \mathbf{M} \models_w \phi \\ \mathbf{M} \models_w \neg\phi_1 \wedge \phi_2 &\text{ iff } \mathbf{M} \models_w \phi_1 \text{ and } \mathbf{M} \models_w \phi_2 \\ \mathbf{M} \models_w \Box_a \phi &\text{ iff for all } v \text{ with } wR_a v: \mathbf{M} \models_v \phi \end{aligned}$$

Action models over this language are defined by:

Definition 4. An **Action Model** is a tuple $(E, \mathbf{P}, \mathbf{pre})$ where

- E is a non-empty set of events.
- \mathbf{P} is a function that assigns to every agent $a \in A$ a binary relation R_a on E .
- **pre** is a precondition function that assigns to every $e \in E$ a formula from \mathcal{L} .

From now on we call the regular epistemic models *static models*.

Updating a static model $\mathbf{M} = (W, R, V)$ with an action model $\mathbf{A} = (E, \mathbf{P}, \mathbf{pre})$ is defined as follows:

Definition 5. The update of static model $\mathbf{M} = (W, R, V)$ with an action model $\mathbf{A} = (E, \mathbf{P}, \mathbf{pre})$ succeeds if the set

$$\{(w, e) \mid w \in W, e \in E, \mathbf{M}, w \models \mathbf{pre}(e)\}$$

is non-empty. The update result is a new static model $\mathbf{M} \otimes \mathbf{A} = (W', R', V')$ with


- $W' = \{(w, e) \mid w \in W, e \in E, \mathbf{M}, w \models \mathbf{pre}(e)\}$,
- R'_a is given by $\{(w, e), (v, f) \mid (w, v) \in R_a, (e, f) \in \mathbf{P}_a\}$,
- $V'(w, e) = V(w)$.

If the static model has a set of distinctive states W_0 and the action model a set of distinctive events E_0 , then the distinctive worlds of $\mathbf{M} \otimes \mathbf{A}$ are the (w, e) with $w \in W_0$ and $e \in E_0$. The distinctive states are the states that can turn out to be the actual state

of a static model. The distinctive events are the events that can turn out to be the actual event of an event model.


Below is an example pair of a static model with an update action model. The static model, on the left, pictures the result of a hidden coin toss, with three onlookers, Alice, Bob and Carol. The model has two distinctive worlds, marked in grey; h in a world means that the valuation makes h true, \bar{h} in a world means that the valuation makes h false in that world. The value h true means that the coin is facing heads up. The fact that both possibilities are distinctive means that both of these could turn out to be the actual world.

The R_a relations for the agents are assumed to be equivalences; reflexive loops for a, b, c at each world are omitted from the picture.



The static model on the left abbreviates two situations: the situation where the coin is facing heads up and it is common knowledge among a, b, c that no-one knows this, and the situation where the coin is showing tails, and again it is common knowledge among a, b, c that no-one knows this. Imagine a situation where one of the agents tosses a coin under a cup and nobody has yet taken a look.

The action model on the right represents a test that reveals to a that the result of the toss is h , while b and c learn that a has learned the answer (without learning the answer themselves). Imagine the act of a taking a peek under the cup, with b and c looking on, but without a revealing to b and c what she sees. The distinctive event of the update is marked grey. The \mathbf{P}_i relations are drawn, for two agents b, c . Reflexive loops are not drawn, so we do not see the \mathbf{P}_i relation for a . The result of the update is shown here:



The result of the update is that the distinction mark on the \bar{h} world has disappeared, that a now knows that the coin is showing heads, that b and c now know that a knows the face of the coin, but that b and c do not know, and all of this is common knowledge. In other words, the model makes each of the following formulas true in its actual world:

$$\begin{aligned}
& \Box_a h, \neg \Box_b h, \neg \Box_c h, \\
& \Box_a(\Box_a h \vee \Box_a \neg h), \Box_b(\Box_a h \vee \Box_a \neg h), \Box_c(\Box_a h \vee \Box_a \neg h), \\
& \Box_a(\neg \Box_b h \wedge \neg \Box_b \neg h), \Box_b(\neg \Box_b h \wedge \neg \Box_b \neg h), \Box_c(\neg \Box_b h \wedge \neg \Box_b \neg h), \\
& \Box_a(\neg \Box_c h \wedge \neg \Box_c \neg h), \Box_b(\neg \Box_c h \wedge \neg \Box_c \neg h), \Box_c(\neg \Box_c h \wedge \neg \Box_c \neg h), \\
& \Box_a \Box_b(\Box_a h \vee \Box_a \neg h), \dots
\end{aligned}$$

The update operator, viewed abstractly, produces a restriction of a product of two models. It is folklore from model theory that a sentence of first order logic is preserved under restriction and product iff the sentence is universal Horn. A universal Horn sentence of FOL is the universal closure of a disjunction with at most one atom disjunct, while the remaining disjuncts are negations of atoms (see, e.g., [52]). The classes of S5 models or S5n models (multi-modal logics where all modalities are S5) are universal Horn: the formulas for reflexivity, symmetry and transitivity can be written as Horn formulas. The classes of KD45 models or KD45n models are not universal Horn, for the seriality requirement cannot be expressed as a universal Horn sentence. Therefore, updating a static model where the accessibilities are KD45 with an update model where the accessibilities are KD45 does not guarantee a result where the accessibilities are again KD45.

6 Logics of Public Announcement

The language of public announcement logic is the extension of \mathcal{L} with an operator $[\phi]\psi$ expressing that after public announcement of ϕ the formula ψ is true in the resulting model.

Definition 6. If $\mathbf{M} = (W, R, V)$ and $\phi \in \mathcal{L}$, then $\mathbf{M}^\phi = (W^\phi, R^\phi, V^\phi)$ is given by:

- $W^\phi = \{w \in W \mid \mathbf{M} \models_w \phi\}$.
- $R^\phi = \lambda a. \{(w, v) \mid w \in W^\phi, v \in W^\phi, wR_a v\}$.
- V^ϕ is the restriction of V to W^ϕ .

Then:

$$\mathbf{M} \models_w [\phi]\psi \text{ iff } \mathbf{M} \models_w \phi \text{ implies } \mathbf{M}^\phi \models_w \psi.$$

The logic of public announcements is now given by the reduction axioms from [59]:

Definition 7. The proof system for public announcement logic consists of the axioms and rules for multimodal S5 epistemic logic (see [19]), plus the following reduction axioms:

$$\begin{aligned}
& \text{Atoms} \vdash [\phi]p \leftrightarrow (\phi \rightarrow p) \\
& \text{Partial functionality} \vdash [\phi]\neg\psi \leftrightarrow (\phi \rightarrow \neg[\phi]\psi) \\
& \text{Distribution} \vdash [\phi](\psi_1 \wedge \psi_2) \leftrightarrow ([\phi]\psi_1 \wedge [\phi]\psi_2) \\
& \text{Knowledge announcement} \vdash [\phi]\Box_a \psi \leftrightarrow (\phi \rightarrow \Box_a[\phi]\psi)
\end{aligned}$$

plus the rules of inference for announcement generalization, given by:

From $\vdash \psi$ infer $\vdash [\phi]\psi$.

These axioms provide a recursive definition of the effect of public announcement, and they can be used to turn every formula from the enhanced language into an equivalent \mathcal{L} formula. This allows us to prove completeness of the logic by means of a reduction argument.

Such an easy complete axiomatization is no longer available when we add an operator for common knowledge $C_B\phi$ to the language, where $C_B\phi$ expresses that ϕ is common knowledge for the agents in $B \subseteq A$. The semantics is given by:

Definition 8. Common knowledge among B :

$$\mathbf{M} \models_w C_B\phi \text{ iff } \mathbf{M} \models_w \phi \text{ for all } v \text{ with } (w, v) \in (R_B)^+,$$

where $R_B = \bigcup_{a \in B} R_a$, and $(R_B)^+$ denotes the transitive closure of R_B .

Still, the reduction method applies, once we are able to express the semantic intuitions for achieving common knowledge by announcement. Introduce an operator for relativized common knowledge, $C_B(\phi, \psi)$, with semantics given by:

Definition 9. Relativized common knowledge among B :

$$\mathbf{M} \models_w C_B(\phi, \psi) \text{ iff } \mathbf{M} \models_w \psi \text{ for all } v \text{ with } (w, v) \in (R_B^\phi)^+,$$

where $R_B^\phi = R_B \cap (W \times \{w \in W \mid \mathbf{M} \models_w \phi\})$.

Intuitively, $C_B(\phi, \psi)$ expresses that every ϕ -path through B accessibilities ends in a state where ψ holds. Let $E_B\phi$ abbreviate $\bigwedge_{a \in B} \Box_a\phi$. Then $E_B\phi$ expresses that ϕ is general knowledge among the agents in $B \subseteq A$.

Definition 10. The proof system for public announcement logic with relativized common knowledge consists of following axioms and rules:

$$\begin{aligned} &\text{Tautologies} && \text{All instances of propositional tautologies} \\ &\text{Knowledge Distribution} && \vdash \Box_a(\phi \rightarrow \psi) \rightarrow (\Box_a\phi \rightarrow \Box_a\psi) \\ &\text{Common Knowledge Distrib} && \vdash C_B(\phi, \psi \rightarrow \chi) \rightarrow (C_B(\phi, \psi) \rightarrow C_B(\phi, \chi)) \\ &\text{Mix} && \vdash C_B(\phi, \psi) \leftrightarrow E_B(\phi \rightarrow (\psi \wedge C_B(\phi, \psi))) \\ &\text{Induction} && \vdash (E_B(\phi \rightarrow \psi) \wedge C_B(\phi, \psi \rightarrow E_B(\phi \rightarrow \psi))) \rightarrow C_B(\phi, \psi) \end{aligned}$$

plus the following inference rules:

$$\begin{aligned} &\text{Modus Ponens} && \text{From } \vdash \phi \text{ and } \vdash \phi \rightarrow \psi \text{ infer } \vdash \psi. \\ &\Box \text{ Necessitation} && \text{From } \vdash \phi \text{ infer } \vdash \Box_a\phi. \\ &C \text{ Necessitation} && \text{From } \vdash \phi \text{ infer } \vdash C_B(\psi, \phi). \end{aligned}$$

The completeness of this system is proved in [16]. To better understand what goes on in the proof system, it is helpful to translate the statements of public announcement with relativized common knowledge into propositional dynamic logic

(PDL), and to note that the above proof system essentially follows the usual PDL axioms. E.g., $C_B(\phi, \psi)$ gets the following PDL translation:

$$[(\bigcup_{a \in B} a; ?\phi)^+] \psi.$$

The connection with PDL will be worked out in the next Section.

7 Connecting up with Epistemic PDL

PDL was designed as a general logic of (computational) action, as a generalization of Floyd-Hoare logic [32, 44]. In Floyd-Hoare logic, one studies correctness statements about programs, such as the following:

$$\begin{array}{l} \{N = \text{gcd}(x, y) \wedge x \neq y\} \\ \quad \text{if } x > y \text{ then } x := x - y \text{ else } y := y - x \\ \{N = \text{gcd}(x, y)\} \end{array}$$

Here the assertion $\{N = \text{gcd}(x, y) \wedge x \neq y\}$ is called the precondition and the assertion $\{N = \text{gcd}(x, y)\}$ the postcondition for the conditional program statement.

The Hoare specification asserts that the loop step in Euclid's GCD algorithm is correct: if x and y store integer numbers that are different, then their GCD does not change if you replace the largest number by the difference of the two numbers.

The general meaning of the Hoare triple $\{\phi\} \pi \{\psi\}$ is: if a state satisfies the precondition ϕ , and program π is executed in that state, then any state that results from this execution will satisfy the postcondition ψ .

Vaughan Pratt saw that such Hoare triples can be viewed as implications in a logic where the program π appears as a modality [60]. The PDL guise of the Hoare correctness statement

$$\{\phi\} \pi \{\psi\}$$

is

$$\phi \rightarrow [\pi] \psi.$$

Hoare logic over programs for integer assignment is undecidable because it talks about variable assignment in the language of first order logic. It has rules like precondition strengthening and postcondition weakening:

$$\frac{\mathbb{N} \models \phi' \rightarrow \phi \quad \{\phi\} \pi \{\psi\}}{\{\phi'\} \pi \{\psi\}}$$

$$\frac{\{\phi\} \pi \{\psi\} \quad \mathbb{N} \models \psi \rightarrow \psi'}{\{\phi\} \pi \{\psi'\}}$$

These use first order statements about natural numbers, which may be undecidable:

$$\mathbb{N} \models \phi' \rightarrow \phi$$

But an extra abstraction step makes the logic decidable again. The basic building blocks for programs in Hoare logic are variable assignment statements $x := E$. Just replace these by arbitrary atomic actions a , and stipulate that the interpretation of a is some binary relation on an abstract set of states. Then PDL emerges as a general program logic, with assertions (formulas) and programs defined by mutual recursion, as follows (assume p ranges over a set of basic propositions $Prop$ and a over a set of basic actions Act):

Definition 11. PDL language:

$$\begin{aligned} \phi &::= \top \mid p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid [\pi]\phi \\ \pi &::= a \mid ?\phi \mid \pi_1; \pi_2 \mid \pi_1 \cup \pi_2 \mid \pi^* \end{aligned}$$

This language is to be interpreted in multimodal Kripke models $\mathbf{M} = (W, R, V)$, where W is a set of worlds or states, R is a function that assigns to every $a \in Act$ a binary relation $R_a \subseteq W^2$, and V is a valuation function that assigns to every $p \in Prop$ a subset of W .

Definition 12. Semantics of PDL. Let $\mathbf{M} = (W, R, V)$. The interpretations $\llbracket \phi \rrbracket^{\mathbf{M}}$ of formulas are subsets of W , and the interpretations $\llbracket \pi \rrbracket^{\mathbf{M}}$ are subsets of W^2 . The clauses for the propositional atoms and the Boolean operators are as usual, the clause for $[\pi]\phi$ is

$$\{w \in W \mid \text{if for all } v \text{ with } w \llbracket \pi \rrbracket^{\mathbf{M}} v : v \in \llbracket \phi \rrbracket^{\mathbf{M}}\},$$

and the clauses for the programs are given by:

$$\begin{aligned} \llbracket a \rrbracket^{\mathbf{M}} &= R_a \\ \llbracket ?\phi \rrbracket^{\mathbf{M}} &= \{(w, w) \in W^2 \mid w \in \llbracket \phi \rrbracket^{\mathbf{M}}\} \\ \llbracket \pi_1; \pi_2 \rrbracket^{\mathbf{M}} &= \llbracket \pi_1 \rrbracket^{\mathbf{M}} \circ \llbracket \pi_2 \rrbracket^{\mathbf{M}} \\ \llbracket \pi_1 \cup \pi_2 \rrbracket^{\mathbf{M}} &= \llbracket \pi_1 \rrbracket^{\mathbf{M}} \cup \llbracket \pi_2 \rrbracket^{\mathbf{M}} \\ \llbracket \pi^* \rrbracket^{\mathbf{M}} &= (\llbracket \pi \rrbracket^{\mathbf{M}})^* \end{aligned}$$

Note the regular operations on relation on the righthand side: \circ for relational composition, \cup for union of relations, and $*$ for Kleene star or reflexive transitive closure. Thus, the complex modalities are handled by the regular operations on relations.

We employ the usual abbreviations: \perp is shorthand for $\neg\top$, $\phi_1 \vee \phi_2$ is shorthand for $\neg(\neg\phi_1 \wedge \neg\phi_2)$, $\phi_1 \rightarrow \phi_2$ is shorthand for $\neg(\phi_1 \wedge \neg\phi_2)$, $\phi_1 \leftrightarrow \phi_2$ is shorthand for $(\phi_1 \rightarrow \phi_2) \wedge (\phi_2 \rightarrow \phi_1)$, $\langle \pi \rangle \phi$ is shorthand for $\neg[\pi]\neg\phi$, and $[\pi^+]\phi$ is shorthand for $[\pi; \pi^*]\phi$.

We now get that $[\pi]\phi$ is true in world w of \mathbf{M} if it holds for all v with $(w, v) \in \llbracket \pi \rrbracket^{\mathbf{M}}$ that ϕ is true in v , and the Hoare assertion $\phi_1 \rightarrow [\pi]\phi_2$ is true in a world w if truth of ϕ_1 in w implies that it holds for all v with $(w, v) \in \llbracket \pi \rrbracket^{\mathbf{M}}$ that ϕ_2 is true in v .

Definition 13. The PDL language is completely axiomatized by the following PDL rules and axioms ([61, 48]):

Modus ponens and axioms for propositional logic
 Modal generalisation From $\vdash \phi$ infer $\vdash [\pi]\phi$

Normality $\vdash [\pi](\phi \rightarrow \psi) \rightarrow ([\pi]\phi \rightarrow [\pi]\psi)$
 Test $\vdash [?\phi]\psi \leftrightarrow (\phi \rightarrow \psi)$
 Sequence $\vdash [\pi_1; \pi_2]\phi \leftrightarrow [\pi_1][\pi_2]\phi$
 Choice $\vdash [\pi_1 \cup \pi_2]\phi \leftrightarrow ([\pi_1]\phi \wedge [\pi_2]\phi)$
 Mix $\vdash [\pi^*]\phi \leftrightarrow (\phi \wedge [\pi][\pi^*]\phi)$
 Induction $\vdash (\phi \wedge [\pi^*](\phi \rightarrow [\pi]\phi)) \rightarrow [\pi^*]\phi$

In the previous Section, we already saw specific instances of the Mix and Induction axioms, in the proof system for public announcement logic with relativized common knowledge.

When the PDL language was designed, the basic actions a were thought of as abstract versions of basic programs (variable assignment statements, say). But nothing in the formal design forces this interpretation. The basic actions could be anything. PDL is a generic action logic for talking about actions as transitions from states of the world to other states of the world.

In PDL, no constraints are imposed on what the actions are. These could be changes in the world, but they could also be epistemic relations. *Epistemic PDL* is just PDL, but with the understanding that the accessibility relations express *knowledge or belief* of agents.

Two extensions of the language are useful: an extension with a global modality G and an extension with a converse operator \checkmark .

Definition 14. Epistemic PDL Action Expressions:

$$\pi ::= a \mid G \mid ?\phi \mid \pi_1; \pi_2 \mid \pi_1 \cup \pi_2 \mid \pi^* \mid \pi^\checkmark$$

Interpretation of the additions:

Definition 15. Semantics of G and \checkmark :

$$\begin{aligned} \llbracket G \rrbracket &= W^2 \\ \llbracket \pi^\checkmark \rrbracket^{\mathbf{M}} &= (\llbracket \pi \rrbracket^{\mathbf{M}})^\checkmark \end{aligned}$$

Thus, $[G]\phi$ expresses that everywhere in the model ϕ holds, and $\langle G \rangle \phi$ expresses that ϕ holds somewhere.

Definition 16. Proof system for epistemic PDL. Axioms and rules of PDL, plus the following. Axioms expressing that G is an S5-operator:

Reflexivity $\vdash \phi \rightarrow \langle G \rangle \phi$
 Symmetry $\vdash \phi \rightarrow [G]\langle G \rangle \phi$
 Transitivity $\vdash \langle G \rangle \langle G \rangle \phi \rightarrow \langle G \rangle \phi$
 Inclusion $\vdash \langle G \rangle \phi \rightarrow \langle \pi \rangle \phi$

Axioms for converse expressing the equivalences that reduce converse over a complex action to converse over atomic actions:

$$\begin{array}{ll}
R_a \subseteq R_a^\smile & \vdash \phi \leftrightarrow [a]\langle a^\smile \rangle \phi \\
R_a^\smile \subseteq R_a & \vdash \phi \leftrightarrow [a^\smile]\langle a \rangle \phi \\
\text{Reduction for } G & \vdash \langle G^\smile \rangle \phi \leftrightarrow \langle G \rangle \phi \\
\text{Reduction for } ?\phi & \vdash \langle ?\phi^\smile \rangle \phi \leftrightarrow \langle ?\phi \rangle \phi \\
\text{Reduction for } ; & \vdash \langle (\pi_1; \pi_2)^\smile \rangle \phi \leftrightarrow \langle \pi_2^\smile; \pi_1^\smile \rangle \phi \\
\text{Reduction for } \cup & \vdash \langle (\pi_1 \cup \pi_2)^\smile \rangle \phi \leftrightarrow \langle \pi_1^\smile \cup \pi_2^\smile \rangle \phi \\
\text{Reduction for } * & \vdash \langle (\pi^*)^\smile \rangle \phi \leftrightarrow \langle (\pi^\smile)^* \rangle \phi
\end{array}$$

It is well-known that adding global modality and converse to PDL does not change its computational properties: the logic remains decidable, satisfiability remains EXPTIME-complete [19, 42], and model checking is PTIME-complete [49]. Many further variations on the set-up of PDL are possible, and we will explore some of them below.

Note that it is not necessary to impose KD45 axioms for belief or S5 axioms for knowledge. Instead, we interpret the atoms as proto epistemic accessibility relations, and we construct the appropriate operators [28]. Let the interpretations of the atomic actions a be an arbitrary binary relation R_a .

Define \sim_a as $(a \cup a^\smile)^*$. This operator is interpreted as the relation

$$(R_a \cup R_a^\smile)^*,$$

and this relation is symmetric, reflexive and transitive. Therefore, \sim_a is an appropriate S5 operator for knowledge.

Some logicians (including Hintikka) have argued for dropping the symmetry requirement, and propose to use an S4 modality for knowledge. The corresponding epistemic PDL operator would be a^* .

To get a KD45 relation from an arbitrary binary relation R_a , consider: $(?[a]\perp; ?\top) \cup a; (a^\smile; a)^*$. Its interpretation S is the relation

$$\{(x, x) \mid \neg \exists z : (x, z) \in R_a\} \cup (R_a \circ (R_a^\smile \circ R_a)^*).$$

Then S is serial, for let x be an arbitrary member of the state set A . If there is no $y \in A$ with $(x, y) \in R_a$ we have $(x, x) \in S$. If there is such a y then $(x, y) \in S$. So in any case there is a $z \in A$ with $(x, z) \in S$.

It is also easy to see that S is transitive and euclidean. Therefore $(?[a]\perp; ?\top) \cup a; (a^\smile; a)^*$ can serve as a KD45 operator, and we have an appropriate way to interpret KD45 belief in epistemic PDL.

Abbreviate this operator as \triangleleft_a . Observe that the interpretation of \triangleleft_a is included in that of \sim_a , so the following principle holds:

$$\langle \triangleleft_a \rangle \phi \rightarrow \langle \sim_a \rangle \phi.$$

Contraposition gives:

$$[\sim_a]\phi \rightarrow [\leq_a]\phi.$$

This expresses that individual knowledge implies individual belief.

Also, the interpretation of $\leq_a; \sim_a$ is included in that of \sim_a , so the following holds as well:

$$[\sim_a]\phi \rightarrow [\leq_a][\sim_a]\phi.$$

Therefore, this interpretation of belief also gives reasonable connections between knowledge and belief in epistemic PDL.

Note that a relation R_a is S5 iff it holds that a and \sim_a have the same interpretation. Similarly, a relation R_a is KD45 iff it holds that a and \leq_a have the same interpretation.

The expressive power of epistemic PDL comes to the fore when we consider common knowledge and shared belief. An appropriate operator for common knowledge between agents a and b is readily defined as

$$(a \cup a^\sim \cup b \cup b^\sim)^*,$$


or equivalently as $(\sim_a \cup \sim_b)^*$. This gives the equivalence relation with every state that is reachable by means of arbitrary numbers of forward or backward a or b steps in a single equivalence class, and this is how common knowledge was explained by philosophers [50], sociologists [33], economists [4] and computer scientists [41].

This is another good reason for using epistemic PDL rather than a logic with explicit modal operators \sim_a and \leq_a . Logics with explicit knowledge and belief modalities are naturally interpreted with respect to models where the corresponding relations have the appropriate properties. But for the case of belief, there is no guarantee that updating such models with ‘reasonable’ update models (where the accessibilities satisfy the same belief properties) results in a new model in the same class. If one uses epistemic PDL, there is no such problem. One can start with a model where the relation of an agent a is KD45 (meaning that a and \leq_a have the same interpretation in that model), and after an update it may turn out that a and \leq_a no longer have the same interpretation. In such a model \leq_a still is a KD45 operator, so there is no problem in interpreting statements about belief.


8 Adding Factual Change

Factual change was already added to update models in LCC, the logic of communication and change of [16], which is basically a system that extends epistemic PDL with generic action update modalities, where the action update can also change the facts of the world. See also [23].


Consider again the model where a has found out the value of the coin, while b and c were onlookers.




Here is a representation of the action of tossing the coin again, with a, b, c present, but without showing the result to any of a, b, c . Explanation: if the coin is tossed again, either the value of h does not change (expressed by \top), or it flips from True to False or vice versa (expressed by $h := \neg h$). Another way of representing this would be by generating a history of coin flips h_1, h_2 , and so on.




After an update with this fact changing action above, we get:



This model looks complicated, but it is bisimilar to the following model, where the coin may have fallen either way, and it is common knowledge that none of the agents knows which side is up. And that's intuitively right, for all players are aware that nobody knows anything about how the coin has fallen this time.



Another example of a simple fact-changing update is the flip of a coin to its other side, represented by the propositional substitution $h := \neg h$. For example, imagine a situation where b and c are aware of the fact that a coin is flipped to its other side (say by some trusted agent d), while a mistakenly believes that nothing has happened. This is modelled by the following action model:



Note that we can no longer omit reflexive arrows, for this model is not reflexive for the a relation.

Consider again the model where a has found out the value of the coin, while b and c were onlookers. After an update with the fact changing action above, we get the result in Figure 5.

Now $[\langle a \rangle][\sim_a]h$ is true in the actual situation: a believes that she knows that the coin is showing heads. But $[(\sim_b \cup \sim_c)^*](\neg[\sim_a]h \wedge \neg[\sim_a]\neg h)$ is also true: b and c have common knowledge that a does *not* know what the coin is showing. Nobody knows what the coin is showing, but b and c know that a is mistaken about what she knows. This kind of situation occurs very often in everyday life: something happens, we mistake it for something else, and we end up with a false belief. This is one of the ways in which knowledge can degrade to mere belief.

9 Adding Belief Change

We will now also add belief change. This was not yet present in LCC, but it is studied extensively in Johan van Benthem's subsequent work [13]. See also [26] and [28]. As a first example, consider the coin situation above, where a has been led astray by failing to observe some change in the world.




Fig. 5 A failure to observe a fact-changing events leading to false belief

Suppose a suddenly comes to believe that she was led astray. She (publicly) updates her belief by accepting that every conceivable state of affairs might be the true state of affairs. The action model for that is:

$$a := \sim a$$

The result of updating the epistemic model of Figure 5 with this is again an S5 model (so we can omit reflexive loops): see Figure 6.

We see that relational change extends the expressive power of the updates, for it can *add* arrows, while action model update without relational change can only *delete* arrows.

We will now formally state how to modify the update process to accommodate both factual change and belief change. Let an action model with both kinds of changes be a quintuple

$$\mathbf{A} = (E, \mathbf{P}, \mathbf{pre}, \mathbf{Sub}, \mathbf{SUB})$$

where $E, \mathbf{P}, \mathbf{pre}$ are as before, \mathbf{Sub} is a function that assigns a propositional binding (or propositional substitution) to each $e \in E$, and \mathbf{SUB} is a function that assigns a




Fig. 6 A belief revision that creates an S5 model from a KD45 model

relational binding to each $e \in E$. A propositional binding is a map from proposition letters to formulas, represented by a finite set of links

$$\{p_1 \mapsto \phi_1, \dots, p_n \mapsto \phi_n\}$$

where the p_k are all different, and where no ϕ_k is equal to p_k . It is assumed that each p that does not occur in a left-hand side of a binding is mapped to itself.

Similarly, a relational binding is a map from agents to program expressions, represented by a finite set of links

$$\{a_1 \mapsto \pi_1, \dots, a_n \mapsto \pi_n\}$$

where the a_j are agents, all different, and where the π_j are program expressions from the PDL language. It is assumed that each a that does not occur in the left-hand side of a binding is mapped to a . Use ε for the identity propositional or relational substitution.

Definition 17. The update execution of static model $\mathbf{M} = (W, P, V)$ with action model $\mathbf{A} = (E, \mathbf{P}, \mathbf{pre}, \mathbf{Sub}, \mathbf{SUB})$ is a tuple

$$\mathbf{M} \otimes \mathbf{A} = (W', P', V')$$

where:

- $W' = \{(w, e) \mid \mathbf{M}, w \models \mathbf{pre}(e)\}$.
- P'_a is given by

$$\{(w_1, e_1), (w_2, e_2) \mid \text{there is a } \mathbf{SUB}(e_1)(a) \text{ path from } (w_1, e_1) \text{ to } (w_2, e_2) \text{ in } \mathbf{M} \otimes \mathbf{A}\}.$$

- $V'(p) = \{(w, e) \in W' \mid \mathbf{M}, w \models \mathbf{Sub}(e)(p)\}$.

The definition of P'_a refers to paths in the old style update product which is denoted with \otimes .

Consider the suggestive upgrade $\#_a\phi$ discussed in Van Benthem and Liu [15] as a relation changer (*uniform* relational substitution):

$$\#_a\phi =_{\text{def}} ?\phi; a; ?\phi \cup ?\neg\phi; a; ?\neg\phi \cup ?\neg\phi; a; ?\phi.$$

This models a kind of belief change where preference links from ϕ worlds to $\neg\phi$ worlds for agent a get deleted. It can be modelled as the following example of public belief change.

Public Belief Change: Action model

$$G = (\{e\}, \mathbf{P}, \mathbf{pre}, \mathbf{Sub}, \mathbf{SUB})$$

where:

- For all the $i \in \text{Ag}$, $\mathbf{P}_i = \{(e, e)\}$.
- $\mathbf{pre}(e) = \top$.
- $\mathbf{Sub}(e) = \varepsilon$.
- $\mathbf{SUB}(e) = \{a \mapsto \#_a\phi, b \mapsto \#_b\phi\}$.

In a picture (reflexive arrows omitted):

$$a := \#_a\phi, b := \#_b\phi$$

Note that our definition of \otimes update implements *point-wise* relational substitutions, which is a more powerful mechanism than merely upgrading the relations *uniformly* everywhere in the model. This is illustrated by the following example.

Non-public Belief Change: Action model

$$G' = (\{e_0, e_1\}, \mathbf{P}, \mathbf{pre}, \mathbf{Sub}, \mathbf{SUB})$$


where:

- For all $i \in \text{Ag}$, if $i \neq b$ then $\mathbf{P}_i = \{(e_0, e_0), (e_1, e_1)\}$,
 $\mathbf{P}_b = \{(e_0, e_0), (e_1, e_1), (e_0, e_1), (e_1, e_0)\}$

- $\mathbf{pre}(e_0) = \mathbf{pre}(e_1) = \top$.
- $\mathbf{Sub}(e_0) = \mathbf{Sub}(e_1) = \varepsilon$.
- $\mathbf{SUB}(e_0) = \{a \mapsto \#_a\phi\}$, $\mathbf{SUB}(e_1) = \varepsilon$.

Assume e_0 is the actual event.

This changes the belief of a while b remains unaware of the change. In a picture (reflexive arrows omitted, since this is an S5 model):



Let PDL^+ be the result of adding modalities of the form $[\mathbf{A}, e]\phi$ to PDL, with the following interpretation clause:

$$\mathbf{M}, w \models [\mathbf{A}, e]\phi \text{ iff } \mathbf{M}, w \models \mathbf{pre}(e) \text{ implies } \mathbf{M} \otimes \mathbf{A}, (w, e) \models \phi.$$

Then the completeness result for LCC extends to a completeness result for PDL^+ . This can be proved by a patch of the LCC completeness proof in [16] where the action modalities are pushed through program modalities by program transformations.

10 Example: Navigation

Navigation problems provide a nice example of the interaction of information flow and change in the world. Consider the case of a robot in a maze. Assume a grid where the robot can move through a sequence of rooms, in some of the four directions North, East, South and West, but some of these directions may be blocked. Assume that the robot has a compass and a map, and that the robot can observe what its present location looks like (which of the four exits is blocked), but not what the next room looks like.

We can assume that the grid and its map look the same. Since the robot has a compass, it knows how to orient the map. Therefore, as soon as the robot uses its sensor it can distinguish the kind of room it is in. There are 15 possibilities:

$$\uparrow, \downarrow, \rightarrow, \leftarrow, \updownarrow, \leftrightarrow, \uparrow\rightarrow, \uparrow\leftarrow, \downarrow\rightarrow, \downarrow\leftarrow, \updownarrow\leftrightarrow, \updownarrow\leftarrow, \up\rightarrow\leftarrow, \up\downarrow\leftarrow.$$

The possibility where the robot can get nowhere is ruled out: we assume the robot is not locked in a room. Initially the robot knows it could be anywhere in the grid, i.e., anywhere on the map. As soon as the robot senses it is in a room of (say) type $\updownarrow\leftrightarrow$, the update that the robot makes is with the observation

$$\langle \uparrow \rangle \top \wedge \langle \rightarrow \rangle \top \wedge \langle \downarrow \rangle \top \wedge \langle \leftarrow \rangle \top.$$

This is to say: I am in a position where I can go North, East, South and West. This rules out all possible locations on the map except for those of type $\updownarrow\leftrightarrow$.

So we assume that the rooms are states in a Kripke model, and that the modalities $\langle \uparrow \rangle$, $\langle \rightarrow \rangle$, $\langle \downarrow \rangle$ and $\langle \leftarrow \rangle$ are available, for moving one step in the indicated directions.

We can use modal formulas to express some obvious constraints on the model. The following formulas fix the relations between the directions:

$$\begin{aligned} & [\uparrow] \langle \downarrow \rangle \top \\ & [\downarrow] \langle \uparrow \rangle \top \\ & [\rightarrow] \langle \leftarrow \rangle \top \\ & [\leftarrow] \langle \rightarrow \rangle \top \end{aligned}$$

Here are two equivalent ways to represent the knowledge of the robot. Either assume that there is a basic proposition loc that is true in precisely one state, and represent the uncertainty of the robot as a set of identical maps, each with the proposition loc at a different place, or assume that there is just a single map, with a basic proposition loc pointing at the actual location of the robot and a basic proposition guess that is true at all places on the map that the robot has not yet ruled out as possibilities for where it might be.

Assume the second representation. Then initialize the value of guess to the set of all states on the map. The first update is when the robot uses its sensors to recognize the type of state. Let ϕ_{fit} be an abbreviation of the following formula:

$$\begin{aligned} \phi_{\text{fit}} := & \bigwedge \{ \langle x \rangle \top \mid x \in \{ \uparrow, \rightarrow, \downarrow, \leftarrow \}, \langle x \rangle \top \text{ is true } \} \\ & \wedge \\ & \bigwedge \{ [x] \perp \mid x \in \{ \uparrow, \rightarrow, \downarrow, \leftarrow \}, \langle x \rangle \top \text{ is false } \}. \end{aligned}$$

So if the robot finds itself in some type of location, and has not learned anything yet, but knows that the map is accurate and well oriented, then it will put guess equal to the set of locations that are of the same type as its current location. So the initial thing that the robot learns is:

$$\text{guess} := \phi_{\text{fit}}$$

Now the robot can make a move, and learn from what it sees in the next location. Making a move is changing the location in the maze. So if the robot moves *North*, this is modelled as:

$$\text{loc} := \langle \downarrow \rangle \text{loc}.$$

Explanation: the old location is now *South* of the new location, so to define the new location in terms of the old, we must ‘look back’. If the robot moves *South*, this is modelled as $\text{loc} := \langle \uparrow \rangle \text{loc}$, if the robot moves *East*, this is modelled as $\text{loc} := \langle \leftarrow \rangle \text{loc}$, if the robot moves *West*, this is modelled as $\text{loc} := \langle \rightarrow \rangle \text{loc}$. This was for modelling the change in the actual world. Now let’s model what the robot learns about its location by observing the new location. After moving South, the robot updates its guess by means of:

$$\text{guess} := \langle \uparrow \rangle \text{guess} \wedge \phi_{\text{fit}}$$

And so on for the other directions.

In the other representation, where the robot maintains a set of maps, updating consists of an update of loc as before, followed by a check $\text{loc} \rightarrow \phi_{\text{fit}}$ that eliminates all maps where the new location does not agree with the new observation given by ϕ_{fit} .

This was for the particular case of a maze, but it is clear that any kind of navigable world can be represented as a Kripke model. An interesting case is a model with non-functional accessibilities. If there are a -labelled arrows in different directions, it means that the robot cannot distinguish between two ‘similar’ actions that result in the robot ending up in different locations. In Kripke semantics, this just means that the following action can be executed in more than one way.

$$\text{loc} := \langle a^\vee \rangle \text{loc}.$$

For further information on the epistemics of navigation we refer to [56] and [65].

11 Epistemic Planning and Protocol Languages

Navigation is a specific example of epistemic planning, for which the DEL framework is well-suited, because one can take atomic planning acts as event model updates. This generalizes the classical approach to planning in the presence of noisy sensors [36]. We give a summary, taking our cues from [20, 51, 1, 3].

A planning domain is a state transition system $\Sigma = (S, A, \gamma)$ where S is a finite or recursively enumerable set of states, A is a finite set of actions, and γ is a partial computable function $S \times A \rightarrow S$ (the state transition function). A planning task can now be viewed as a triple (Σ, s_0, G) , where $\Sigma = (S, A, \gamma)$ is a planning domain, s_0 is a state in S , and G is a subset of S (the set of goal states). A solution to a planning task is a finite sequence of actions a_1, a_2, \dots, a_n (a plan), such that $\bar{\gamma}(a_1, a_2, \dots, a_n)$ is defined and $\in G$, where $\bar{\gamma}: A^* \rightarrow S$ is defined by $\bar{\gamma}(a) = \gamma(s_0, a)$, $\bar{\gamma}(\bar{a}, a) = \gamma(\bar{\gamma}(\bar{a}), a)$ if $\bar{\gamma}(\bar{a})$ is defined, undefined otherwise. Informally, a plan succeeds if using the state transition function γ starting from (s_0, a_1) and following the plan, one can reach a goal state. An epistemic planning task is a special case of this where the s_0 is an epistemic state, the set A is a set of finite action models, and the set of goal states is represented by an epistemic formula ϕ_g .

An example of an epistemic plan for letting both a and b know that p without revealing to a that b knows p and without revealing to b that a knows p is first privately communicating p to a and then privately communicating p to b .

It is proved in [20] that the plan existence problem for three agent epistemic planning with factual change is undecidable. This result is strengthened in [3]: even without factual change, the plan existence problem for two agent epistemic planning in an S5 setting is undecidable.

For an intensional view of planning, we can use a version of PDL over action models, to define plan protocols (for still another guise of PDL, as a multi-agent strategy logic, see [27]).

Definition 18. The DEL protocol language for Ag and $Prop$ is given by:

$$\begin{aligned}\phi &::= p \mid \neg\phi \mid \phi \wedge \phi \mid \Box_a\phi \mid [\pi]\phi \\ \pi &::= (\mathbf{A}, e) \mid \pi \cup \pi \mid \pi; \pi \mid \pi^*\end{aligned}$$

where p ranges over $Prop$, a ranges over Ag , and (\mathbf{A}, e) is an action model (without factual change or belief change) for \mathcal{L} with distinguished event e .

The truth conditions for the protocols π are given by:

Definition 19.

$$\begin{aligned}\mathbf{M} \models_w [\mathbf{A}, e]\phi &\text{ iff } \mathbf{M} \models_w \mathbf{pre}(e) \text{ implies } \mathbf{M} \otimes \mathbf{A} \models_{(w,e)} \phi \\ \mathbf{M} \models_w [\pi_1 \cup \pi_2]\phi &\text{ iff } \mathbf{M} \models_w [\pi_1]\phi \text{ and } \mathbf{M} \models_w [\pi_2]\phi \\ \mathbf{M} \models_w [\pi_1; \pi_2]\phi &\text{ iff } \mathbf{M} \models_w [\pi_1][\pi_2]\phi \\ \mathbf{M} \models_w [\pi^*]\phi &\text{ iff for any finite sequence } \pi; \dots; \pi \quad \mathbf{M} \models_w [\pi; \dots; \pi]\phi\end{aligned}$$

It follows from results in [53] that the satisfiability problem for this protocol language is undecidable. It is proved in [3] that the model checking problem for DEL protocols is also undecidable. This follows from a reduction of the plan existence problem to the model checking problem: an epistemic planning task $((\mathbf{M}, w), A, \phi_g)$ has a solution iff $\mathbf{M} \models_w \neg[A^*]\neg\phi_g$ holds.

12 Further Connections

An intriguing question about action model update that so far has only received a partial answer is: When are two action updates the same? More precisely, let us say that action models \mathbf{A} and \mathbf{B} are *equivalent* iff it holds for all static models \mathbf{M} that $\mathbf{M} \otimes \mathbf{A} \Leftrightarrow \mathbf{M} \otimes \mathbf{B}$, where \Leftrightarrow expresses the existence of a bisimulation that connects each distinctive point from $\mathbf{M} \otimes \mathbf{A}$ with a distinctive point from $\mathbf{M} \otimes \mathbf{B}$.

It turns out that action model bisimulation is not the appropriate structural notion to cover equivalence. In [24] a notion of *parametrized action emulation* is defined that characterizes action model equivalence. In [62, 63] this is replaced by a non-parametrized action emulation relation, and it is shown that this characterizes action model equivalence for canonical action models (action models with maximal consistent subsets of an appropriate closure language as preconditions). The question whether non-parametrized action emulation also characterizes action model equivalence for arbitrary action models is still open.

Another intriguing issue is the proper connection between epistemic/doxastic updating and probability theory. Useful overviews of logics of uncertainty are [57] and

[40]. For the connection with DEL, see [29, 47, 46, 14] and the contribution of Kooi and Demey in the present volume. These proposals do not equate knowledge with certainty, but [25] does; this paper proposes a DEL logic (together with an epistemic model checking program) where the following principles hold ($P_a\phi$ is the probability that agent a assigns to ϕ):

Certainty implies Truth $P_a\phi = 1 \rightarrow \phi$.

Positive Introspection into Certainty $P_a\phi = 1 \rightarrow P_a(P_a\phi = 1) = 1$.

Negative Introspection into Certainty $P_a\phi < 1 \rightarrow P_a(P_a\phi < 1) = 1$.

All these probabilistic versions of DEL incorporate Bayesian updating/learning; the difference is in whether Bayesian updates get analyzed as belief revision or as knowledge growth.

Acknowledgements

Thanks to Guillaume Aucher, Alexandru Baltag, Johan van Benthem, Sonja Smets, for illuminating discussion, feedback and encouragement.

References

1. Mikkel Birkegaard Andersen, Thomas Bolander, and Martin Holm Jensen. Conditional epistemic planning. In *Proceedings of JELIA 2012*, volume 7519 of *Lecture Notes in Artificial Intelligence*, pages 94–106, 2012.
2. Guillaume Aucher. A combined system for update logic and belief revision. Master’s thesis, ILLC, Amsterdam, 2003.
3. Guillaume Aucher and Thomas Bolander. Undecidability in epistemic planning. In *Proceedings of IJCAI 2013 (23rd International Joint Conference on Artificial Intelligence)*, August 2013.
4. R.J. Aumann. Agreeing to disagree. *Annals of Statistics*, 4(6):1236–1239, 1976.
5. F. Bacchus, J. Halpern, and H. Levesque. Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence*, 111(1-2):171–208, 1999.
6. A. Baltag and L.S. Moss. Logics for epistemic programs. *Synthese*, 139(2):165–224, 2004.
7. A. Baltag, L.S. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. In I. Bilboa, editor, *Proceedings of TARK’98*, pages 43–56, 1998.
8. A. Baltag, L.S. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. Technical report, Dept of Cognitive Science, Indiana University and Dept of Computing, Oxford University, Bloomington and Oxford, 2003. Updated version of [7].
9. Alexandru Baltag and Sonja Smets. Dynamic belief revision over multi-agent plausibility models. In W. van der Hoek and M. Wooldridge, editors, *Proceedings of LOFT’06*, pages 11–24, Liverpool, 2006. University of Liverpool.
10. Alexandru Baltag and Sonja Smets. The logic of conditional doxastic actions. In Krzysztof Apt and Robert van Rooij, editors, *New Perspectives on Games and Interaction*, volume 5 of *Texts in Logic and Games*, pages 9–31. Amsterdam University Press, Amsterdam, 2008.

11. Alexandru Baltag and Sonja Smets. Group belief dynamics under iterated revision: fixed points and cycles of joint upgrades. In *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*, TARK '09, pages 41–50, New York, NY, USA, 2009. ACM.
12. P. Battigalli and G. Bonanno. Recent results on belief, knowledge and the epistemic foundations of game theory. *Research in Economics*, 53:149–225, 1999.
13. J. van Benthem. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, 2:129–155, 2007.
14. J. van Benthem, J. Gerbrandy, and B. Kooi. Dynamic update with probabilities. *Studia Logica*, 93:67–96, 2009.
15. J. van Benthem and F. Liu. Dynamic logic of preference upgrade. *Journal of Applied Non-Classical Logics*, 14(2):157–182, 2007.
16. J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. *Information and Computation*, 204(11):1620–1662, 2006.
17. Johan van Benthem. Information update as relativization. Technical report, ILLC, Amsterdam, 2000. Available from <http://staff.science.uva.nl/~johan/Upd=Rel.pdf>.
18. Johan van Benthem, Hans van Ditmarsch, Jan van Eijck, and Jan Jaspars. *Logic in Action*. Internet, 2012. electronic book, available from www.logicinaction.org.
19. P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001.
20. Thomas Bolander and Mikkel Birkegaard Andersen. Epistemic planning for single- and multi-agent systems. *Journal of Applied Non-Classical Logics*, 21(1):9–34, 2011.
21. R.T. Cox. Probability, frequency, and reasonable expectation. *American Journal of Physics*, 14(1):1–13, 1946.
22. H.P. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer, 2006.
23. H.P. van Ditmarsch, W. van der Hoek, and B.P. Kooi. Dynamic epistemic logic with assignment. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 05)*, pages 141–148, New York, 2005. ACM Inc.
24. J. van Eijck, J. Ruan, and T. Sadzik. Action emulation. *Synthese*, 185:131–151, 2012.
25. Jan van Eijck. Learning about probability. available from homepages.cwi.nl/~jve/software/prodemo.
26. Jan van Eijck. Yet more modal logics of preference change and belief revision. In K.R. Apt and R. van Rooij, editors, *New Perspectives on Games and Interaction*, volume 4 of *Texts in Logic and Games*, pages 81–104. Amsterdam University Press, 2008.
27. Jan van Eijck. PDL as a multi-agent strategy logic. In Burkhard C. Schipper, editor, *TARK 2013 – Theoretical Aspects of Reasoning About Knowledge, Proceedings of the 14th Conference – Chennai, India*, pages 206–215, 2013.
28. Jan van Eijck and Yanjing Wang. Propositional Dynamic Logic as a logic of belief revision. In Wilfrid Hodges and Ruy de Queiroz, editors, *Proceedings of Wollic'08*, number 5110 in *Lecture Notes in Artificial Intelligence*, pages 136–148. Springer, 2008.
29. R. Fagin and J.Y. Halpern. Reasoning about knowledge and probability. *Journal of the ACM*, pages 340–367, 1994.
30. R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.
31. Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. Knowledge-based programs. *Distrib. Comput.*, 10(4):199–225, 1997.
32. R.W. Floyd. Assigning meanings to programs. In *Proceedings AMS Symposium Applied Mathematics*, volume 19, pages 19–31, Providence, R.I., 1967. American Mathematical Society.
33. M.F. Friedell. On the structure of shared awareness. *Behavioral Science*, 14(1):28–39, 1969.
34. P. Gärdenfors. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. MIT Press, Cambridge Mass, 1988.
35. J. Gerbrandy. *Bisimulations on Planet Kripke*. PhD thesis, ILLC, Amsterdam, 1999.

36. Malik Ghallab, Dana S. Nau, and Paulo Traverso. *Automated Planning: Theory and Practice*. Morgan Kaufmann, 2004.
37. Patrick Girard, Jeremy Seligman, and Fenrong Liu. General dynamic dynamic logic. In Thomas Bolander, Torben Braüner, Silvio Ghilardi, and Lawrence S. Moss, editors, *Advances in Modal Logic*, pages 239–260. College Publications, 2012.
38. W. Groeneveld. *Logical Investigations into Dynamic Semantics*. PhD thesis, ILLC, Amsterdam, 1995.
39. J. Halpern. Using reasoning about knowledge to analyse distributed systems. *Annual Review of Computer Science*, 2:37–68, 1987.
40. J. Halpern. *Reasoning About Uncertainty*. MIT Press, 2003.
41. J.Y. Halpern and Y.O. Moses. Knowledge and common knowledge in a distributed environment. In *Proceedings 3rd ACVM Symposium on Distributed Computing*, pages 50–68, 1984.
42. Edith Hemaspaandra. The price of universality. *Notre Dame Journal of Formal Logic*, 37(2):174–203, 1990.
43. J. Hintikka. *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Cornell University Press, Ithaca N.Y., 1962.
44. C.A.R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):567–580, 583, 1969.
45. J. Jaspars. *Calculi for Constructive Communication*. PhD thesis, ITK, Tilburg and ILLC, Amsterdam, 1994.
46. Barteld P. Kooi. *Knowledge, Chance, and Change*. PhD thesis, Groningen University, 2003.
47. Barteld P. Kooi. Probabilistic dynamic epistemic logic. *Journal of Logic, Language and Information*, 12(4):381–408, 2003.
48. D. Kozen and R. Parikh. An elementary proof of the completeness of PDL. *Theoretical Computer Science*, 14:113–118, 1981.
49. Martin Lange. Model checking propositional dynamic logic with all extras. *Journal of Applied Logic*, 4(1):39–49, March 2006.
50. D.K. Lewis. *Convention: A Philosophical Study*. Harvard University Press, 1969.
51. Benedikt Löwe, Eric Pacuit, and Andreas Witzel. Del planning and some tractable cases. In *Proceedings of the Third international conference on Logic, rationality, and interaction, LORI'11*, pages 179–192, Berlin, Heidelberg, 2011. Springer-Verlag.
52. George McNulty. Fragments of first order logic, i: Universal Horn logic. *Journal of Symbolic Logic*, 1977.
53. J.S. Miller and L.S. Moss. The undecidability of iterated modal relativization. *Studia Logica*, 79(3):373–407, 2005.
54. R.C. Moore. A formal theory of knowledge and action. In J.R. Hobbs and R.C. Moore, editors, *Formal Theories of the Commonsense World*, pages 319–358. Ablex Publishing, Norwood, N.J., 1985.
55. M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
56. Prakash Panangaden and Mehrnoosh Sadrzadeh. Learning in a changing world, an algebraic modal logical approach. In *Proceedings of the 13th international conference on Algebraic methodology and software technology, AMAST'10*, pages 128–141, Berlin, Heidelberg, 2011. Springer-Verlag.
57. Jeff B. Paris. *The uncertain reasoner's companion — a mathematical perspective*, volume 39 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1994.
58. Andrés Perea. *Epistemic Game Theory: Reasoning and Choice*. Cambridge University Press, 2012.
59. J. A. Plaza. Logics of public communications. In M. L. Emrich, M. S. Pfeifer, M. Hadzikadic, and Z. W. Ras, editors, *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems*, pages 201–216, 1989.
60. V. Pratt. Semantical considerations on Floyd–Hoare logic. *Proceedings 17th IEEE Symposium on Foundations of Computer Science*, pages 109–121, 1976.
61. K. Segerberg. A completeness theorem in the modal logic of programs. In T. Traczyck, editor, *Universal Algebra and Applications*, pages 36–46. Polish Science Publications, 1982.

62. Floor Sietsma. *Logics of Communication and Knowledge*. PhD thesis, ILLC, Amsterdam, December 2012.
63. Floor Sietsma and Jan van Eijck. Action emulation between canonical models. *Journal of Philosophical Logic*, Accepted for Publication, 2013.
64. Yanjing Wang. *Epistemic Modelling and Protocol Dynamics*. PhD thesis, ILLC, Amsterdam, 2010.
65. Yanjing Wang and Yanjun Li. Not all those who wander are lost: dynamic epistemic reasoning in navigation. In *Proceedings of AIML 2012*, Copenhagen, 2012. to appear.