

Arrow Update Synthesis

Hans van Ditmarsch^{*,†} Wiebe van der Hoek[‡] Barteld Kooi[§]
Louwe B. Kuijer[¶]

Abstract

In this contribution we present *arbitrary arrow update model logic* (AAUML). This is a *dynamic epistemic logic* or *update logic*. In update logics, static/basic modalities are interpreted on a given relational model whereas dynamic/update modalities induce transformations (updates) of relational models. In AAUML the update modalities formalize the execution of *arrow update models*, and there is also a modality for quantification over arrow update models. Arrow update models are an alternative to the well-known action models. We provide an *axiomatization* of AAUML. The axiomatization is a rewrite system allowing to eliminate arrow update modalities from any given formula, while preserving truth. Thus, AAUML is decidable and equally expressive as the base multi-agent modal logic. Our main result is to establish *arrow update synthesis*: if there is an arrow update model after which φ , we can *construct* (synthesize) that model from φ . We also point out some pregnant differences in *update expressivity* between arrow update logics, action model logics, and refinement modal logic.

Keywords: modal logic, synthesis, dynamic epistemic logic, expressivity

1 Introduction

Modal logic In modal logic we formalize that propositions may not be merely true or false, but that they are necessarily or possibly true or false, or that they may be desirable, or forbidden, or true later, or never, or that they are *known*. A common setting is for such modal propositions to be interpreted in *relational models*, also known as *Kripke models*. They consist of a domain of abstract objects, called states or worlds; then, given a set of labels, often representing agents, for each such agent a binary relation between those

*LORIA, CNRS, France, hans.van-ditmarsch@loria.fr

†Hans van Ditmarsch is corresponding author. He is also affiliated to IMSc, Chennai, India. Support from ERC project EPS 313360 is kindly acknowledged.

‡Computer Science, University of Liverpool, United Kingdom, Wiebe.Van-Der-Hoek@liverpool.ac.uk


§Department of Philosophy, University of Groningen, the Netherlands, b.p.kooi@rug.nl

¶Computer Science, University of Liverpool, United Kingdom, Louwe.Kuijer@liverpool.ac.uk

states; and, finally, a valuation of atomic propositions on the domain, typically seen as a unary relation, i.e., a property satisfied on a subset of the domain. The truth of a modal proposition is relative to a state in the relational model, called the actual state or the *point* of the model. The unit of interpretation is thus a *pointed model*: a pair consisting of a relational model and an actual state.

If a pair (s, s') is in the relation for a this can mean that after executing action a in state s the resulting state is s' . But it can also mean that agent a considers state s' desirable in case she is in state s . The interpretation that we focus on, is that of *information*. That is, it is consistent with a 's information in state s that the state would be s' . In state s it is true that agent a *knows* φ (or *believes* φ , depending on the properties of the relation), notation $\Box_a\varphi$, if the formula φ is true in all states s' accessible from s , i.e., for all s' with (s, s') in the relation for a . The modal logics using that kind of interpretation of modalities are called epistemic logics [7, 28].

As an example, consider two agents a, b (commonly known to be) uncertain about the truth of a propositional variable p . The uncertainty of a and b can be pictured as follows. We ‘name’ the states with the value of the variable p . The actual state is framed. Pairs in the accessibility relation are visualized as labelled *arrows*. In the actual state: p is true, agent a does not know p because she considers a state possible wherein p is false (formally $\neg\Box_a p$), agent a also does not know $\neg p$ because she considers a state possible wherein p is true (formally $\neg\Box_a\neg p$, also written as $\Diamond_a p$), and similarly for agent b . Agent a also knows that she is ignorant about p , as this is true in both states that she considers possible. The accessibility relations for a and b are both equivalence relations. This is always the case if the modalities represent knowledge.



Update logic In this work we focus on modal logics that are *update logics*. Apart from the modalities that are interpreted *in* a relational model, they have other modalities that are interpreted by *transforming* a relational model (and by then interpreting the formulas bound by that modality in the transformed model). If the modal logic is an epistemic logic, update logics are called *dynamic epistemic logics*. To distinguish them we call the former *static* and we call the latter *dynamic*.

The *updates* X that we consider can be defined as transformers of relational models. This transformation induces a binary update relation between pointed models. To an update relation corresponds an *update modality* (often also called update) that is interpreted with this relation, so we can see those as $[X]$ or $\langle X \rangle$, where $[X]\varphi$ means that φ is true in all pointed models transformed according to the X relation, and $\langle X \rangle\varphi$ that there is a pair of pointed models in the relation. Given a relational model we can change its domain of states, the relations between the states, or the valuations of atomic propositions, or two or more of those at the same time. There are therefore many options for change. Change

the valuation of a model is also known as *factual change* [30, 25]. Update involving factual change is an interesting topic, but it is outside the scope of the current paper.

Public announcement logic The basic update for states is the model restriction, and the basic update operation interpreted as a model restriction is a *public announcement*. The logic with epistemic modalities and public announcements is *public announcement logic* (PAL) [22, 6]. A public announcement of φ restricts the domain to all states where the announced formula φ is true, thereby decreasing the uncertainty of the agents. As a result of the domain restriction, the relations and the valuation are adjusted in the obvious way. A condition for the transformation is that the actual state is in the restriction. This means that the announcement formula is true when announced.

As an example, after the public announcement of p , both a and b know that p :

$$\begin{array}{ccc} ab & & ab \\ \downarrow & & \downarrow \\ \neg p & \xleftarrow{ab} & \boxed{p} \end{array} \Rightarrow \begin{array}{ccc} ab & & \\ \downarrow & & \\ & \xleftarrow{ab} & \boxed{p} \end{array}$$

Arrow update logic The basic update for relations is the relational restriction, i.e., a restriction of the *arrows*: a pair in the relation is called an ‘arrow’. This leaves all states intact, although some may have become unreachable. In *arrow update logic* (AUL), proposed in [20] we specify which arrows we wish to preserve, by way of specifying what formulas should be satisfied at the source (state) of the arrow and the target (state) of the arrow. This determines the model transformation. Such a specification is called an *arrow update*. The logic AUL contains modalities for arrow updates.

Given initial uncertainty about p with both agents, a typical arrow update is the action wherein Anne (a) opens an envelope containing the truth about p while Bill (b) observes Anne reading the contents of the letter. We preserve all arrows satisfying one of $p \rightarrow_a p$, $\neg p \rightarrow_a \neg p$, and $\top \rightarrow_b \top$. Therefore, only two arrows disappear, $\neg p \rightarrow_a p$ and $p \rightarrow_a \neg p$.

$$\begin{array}{ccc} ab & & ab \\ \downarrow & & \downarrow \\ \neg p & \xleftarrow{ab} & \boxed{p} \end{array} \Rightarrow \begin{array}{ccc} ab & & ab \\ \downarrow & & \downarrow \\ \neg p & \xleftarrow{b} & \boxed{p} \end{array}$$


The boundary between state elimination and arrow elimination is subtle. If p is true, the following arrow update with $\top \rightarrow_a p$, $\top \rightarrow_b p$ is the same update as a public announcement of p . This is because there is no arrow from the p state to the $\neg p$ state after the update. Therefore, if p is true, the $\neg p$ state does not matter. In another formalism this arrow update is known as the arrow elimination semantics of public announcement [13, 19].

$$\begin{array}{ccc} ab & & ab \\ \downarrow & & \downarrow \\ \neg p & \xleftarrow{ab} & \boxed{p} \end{array} \Rightarrow \begin{array}{ccc} & & ab \\ & & \downarrow \\ \neg p & \xleftarrow{b} & \boxed{p} \end{array}$$

Generalizations In PAL and AUL the complexity (the number of states) of the relational model cannot increase. By generalizing the mechanism underlying state elimination and arrow elimination we can achieve that, and thus express more model transformations. This increases their *update expressivity*. From the perspective of information change, this adds uncertainty about what is happening. We obtain *action models* [6] as a generalization of public announcements, and *arrow update models* [21] as a generalization of arrow updates.

Action model logic *Action model logic* (AML) was proposed by Baltag, Moss and Solecki in [6]. An *action model* is like a relational model but the elements of the domain are called *actions* instead of states, and instead of a valuation a *precondition* is assigned to each domain element. The transformed relational model is then the modal product of the relational model and the action model, restricted to (state,action) pairs where the action can be executed in that state. We refer to Section 6 for a formal introduction.

An example is the action as above wherein Anne reads the contents of a letter containing p or $\neg p$, but now with the increasing uncertainty that Bill is uncertain whether Anne has read the letter (and that they are both aware of these circumstances). The action model is not depicted (details are in Section 6). The model transformation is as follows. In the resulting framed state, a knows that p , but b considers it possible that a is uncertain about p , i.e., $\Box_a p \wedge \Diamond_b \neg(\Box_a p \vee \Box_a \neg p)$. In the figure we assume transitivity of the relation for b .



Similar logics (or semantics) for action composition are found in [26, 18, 25, 2, 36]. Action model logic is often referred to as (the) dynamic epistemic logic. As said, we use the latter more generally, namely to denote any update logic with an epistemic interpretation.

Arrow update model logic *Generalized arrow update logic* [21] is a (indeed) generalization of arrow update logic where the dynamic modalities for information change formalize execution of (pointed) *arrow update models*, structures akin to the *action models* of action model logic. In this contribution, instead of generalized arrow update logic we call it *arrow update model logic* (AUML). The arrow updates of [20] correspond to singleton arrow update models. The next Section 2 formally introduces them. The above is also an example of arrow update model execution — Section 6 explains how to get action models from arrow update models and vice versa, and to what extent they define the same update.

Quantification over information change Another extension of update logics is with quantification over updates. *Arbitrary public announcement logic* (APAL) adds quantification over public announcements to PAL [5]. *Arbitrary arrow update logic* (AAUL) [32] extends arrow update logic with quantifiers over information change induced by arrow updates: it contains dynamic modalities formalizing that *there is an arrow update after which* φ . *Arbitrary action model logic* (AAML) by Hales [14] add quantifiers over action models to AML. In *arbitrary arrow update model logic* (AAUML), the topic of this paper, we add quantifiers over arrow update models to the logical language. It is like Hales’ arbitrary action model logic. *Refinement modal logic* (RML) [9] has a modality representing quantification over updates, but does not have (deterministic/concrete) update modalities in the object language to quantify over. We show that the AAML and AAUML quantifier behave much (but not quite) like the refinement quantifier in RML. Section 7 is devoted to it.

Figure 1 gives an overview of the different logics discussed in the paper, in their relation to AAUML. The four logics in the left square are based on state manipulation, the four logics in the right square are based on arrow manipulation. Entirely on the left we find the base modal logic ML and the logic RML, that is also arrow manipulating.

All these logics are equally expressive as ML and are decidable, which can be shown by truth-preserving rewriting procedures to eliminate updates (for AAUML this is one of the results of the paper), except for APAL and AAUL, which are more expressive and undecidable [5, 11, 32, 33]. However, the logics greatly differ in update expressivity, as the typical examples above already demonstrated. See also Sections 5 — 7. Finally, it should be mentioned that all the logics are invariant under bisimulation. This is because the parameters of the model transforming dynamic modalities and quantifiers are (model restrictions induced by) formulas.

There are many other updates and update logics that we do not consider in this paper. In particular we do not consider updates X that can only be defined as *pointed* model transformers (that is, they cannot be *globally* defined on the entire model; they are defined *locally*: how they transform the model depends on the actual state). If such were the definition of an update, even the interpretation of a static modality can be seen as an update, namely transforming the model with point s into the model with point s' , where the point has shifted given a pair (s, s') in the relation for an agent. Such local update logics are often more expressive than modal logic, are often undecidable, are typically not invariant under (standard) bisimulation, and may lack axiomatizations. Examples are [24, 4, 1]. In [1] not only relational *restriction* is considered but also relational *expansion* (‘bridge’) and relational *change* that is neither restriction nor expansion, such as reversing the direction of arrows (‘swap’). It should finally be noted that the distinction between static modalities, interpreted in a model, and dynamic modalities, interpreted as updates, is not rigid: unifying perspectives include [25].

Synthesis For these update logics we can ask whether there is an update that achieves a certain goal. For the logics without quantification this question cannot be asked in the




Figure 1: An overview of update logics discussed in the paper. Horizontal arrows informally represent *more complex* updates. Vertical arrows informally represent *quantification over* updates. The arrows can be interpreted as syntactic extensions (modulo the names of quantifiers) or as semantic generalizations. Assume transitive closure.

object language but only meta-logically. That is, we can ask whether there is an update X such that $\langle X \rangle \varphi$ is true. For the update logics with quantification this question can be asked in the object language. Let $\langle ? \rangle$ be (the existential version of) that quantifier. Then $\langle ? \rangle \varphi$ asks whether there is an update X that makes φ true.

Only knowing *whether* there is an update that achieves a goal is not very satisfying; we would also like to know *which* update, if any, achieves the goal. So we would like to know not only whether the goal is achievable but also how it can be achieved. The process of constructing this update that achieves the goal is known as *synthesis*.


Formally, the synthesis problem for a given type of update takes as input a formula φ , and gives as output an update X of that type such that, whenever φ can be achieved, then X achieves φ . In symbols, this is the validity of $\langle ? \rangle \varphi \rightarrow \langle X \rangle \varphi$.

This is a rather strong goal. We do not consider it sufficient to find, for every pointed model (M, s) , an update $X_{(M,s)}$ such that (M, s) satisfies $\langle ? \rangle \varphi \rightarrow \langle X_{(M,s)} \rangle \varphi$. We want one single update X_φ that achieves φ in every model where φ is achievable. Because this goal is so strong, there is, in general, no guarantee that synthesis is possible.

For PAL this strong kind of synthesis is impossible. If (M_1, s_1) satisfies $\langle \psi_1 \rangle \varphi$ and (M_2, s_2) satisfies $\langle \psi_2 \rangle \varphi$, so if φ can be achieved in two different situations using two different public announcements, then there is typically no unifying public announcement ψ such that (M_1, s_2) satisfies $\langle \psi \rangle \varphi$ and (M_2, s_2) satisfies $\langle \psi \rangle \varphi$.¹

For AUL this strong kind of synthesis is also not possible. But, somewhat surprisingly,

¹For example, consider the four-state model below; \bar{p} means that p is false in that state, etc. Both states where p, q, r are all true satisfy that $\langle ? \rangle (\Box_a p \wedge \neg \Box_b p)$. In the top-left pqr -state this is true because $\langle q \rangle (\Box_a p \wedge \neg \Box_b p)$ is true, whereas in the bottom-right pqr -state this is true because $\langle r \rangle (\Box_a p \wedge \neg \Box_b p)$ is true. However, there is no announcement φ such that $\langle \varphi \rangle (\Box_a p \wedge \neg \Box_b p)$ is truth in both pqr -states. Assuming that there were such an announcement easily leads to a contradiction.



in [14], Hales showed that this synthesis is possible for AML. This result was surprising for the following reason. Hales obtained his synthesis result with refinement modalities as quantifiers. It was already known that finite action model execution results in a refinement of the current relational model, but also that the other direction does not hold: there are refinements (i.e., updates) that can only be achieved by executing an infinite action model [29]. However, as the synthesis is with respect to making a given formula φ true, a finite syntactic object, synthesis for AML was after all possible.

In this contribution we show that synthesis is also possible for AUML. That is, for a given goal formula φ , we can construct an arrow update model X such that

For all (M, s) : there is an arrow update model Y such that (M, s) satisfies $\langle Y \rangle \varphi$, if and only if (M, s) satisfies $\langle X \rangle \varphi$.

In AAUML we also have a quantifier over arrow update models. Therefore, in that logic the synthesis translates to the above-mentioned validity $\langle ? \rangle \varphi \rightarrow \langle X \rangle \varphi$. In AUML / AAUML we synthesize a (single-)pointed arrow update model, whereas for AAUML Hales synthesizes a *multi*-pointed action model, and it can be easily shown that this cannot be single-pointed.

Results in the paper In this contribution we present *arbitrary arrow update model logic* (AAUML), that further extends arrow update model logic AUML, namely with dynamic modalities formalizing that *there is an arrow update model after which φ* . For this logic AAUML we obtain various results. We provide an *axiomatization* of AAUML. The axiomatization is a rewrite system allowing to eliminate dynamic modalities from any given formula, while preserving truth. Thus, unlike AAUL, AAUML is decidable, and equally expressive as multi-agent modal logic. We establish *arrow update model synthesis*: if there is an arrow update model after which φ , we can *construct* (synthesize) that model from φ . We define a notion of *update expressivity* and we determine the relative update expressivity of AAUML and other arrow update logics and action model logics, and RML.

Overview of content Section 2 presents the syntax and semantics of arbitrary arrow update model logic, AAUML, and elementary structural notions. In Section 3 we describe the procedure for synthesizing arrow update models. In that section we also introduce a number of validities that are useful when introducing an axiomatization for AAUML, which we do in the subsequent Section 4. Section 5 introduces the notion of *update expressivity*. Section 6 compares AAUML and AAML, and in particular their update expressivity. This comparison also includes examples of arrow update models that have exponentially larger corresponding action models. Section 7 compares AAUML to RML.

2 Arbitrary arrow update model logic

Throughout this contribution, let A be a finite set of *agents* and let P be a disjoint countably infinite set of *propositional variables* (or *atoms*).

2.1 Structures

A *relational model* is a triple $M = (S, R, V)$ with S a non-empty *domain* (set) of *states* (also denoted $\mathcal{D}(M)$), R a function assigning to each agent $a \in A$ an *accessibility relation* $R_a \subseteq S \times S$, and $V : P \rightarrow S$ a *valuation function* assigning to each propositional variable $p \in P$ the subset $V(p) \subseteq S$ where the variable is true. For $s \in S$, the pair (M, s) is called a *pointed relational model*, and for $T \subseteq S$, the pair (M, T) is called a *multi-pointed relational model*.

For any relation R on a domain X , instead of $(x, y) \in R$ (where $x, y \in X$) we may write $R(x, y)$ or xRy , and $R(x)$ or Rx for the set $\{y \in X \mid R(x, y)\}$. If $R(x, y)$ we also say that R *links* x to y , or that there *is an arrow* from x to y . Relation R is: *reflexive* iff for all $x \in X$, $R(x, x)$; *serial* iff for all $x \in X$ there is $y \in X$ such that $R(x, y)$; *transitive* iff for all $x, y, z \in X$, if $R(x, y)$ and $R(y, z)$ then $R(x, z)$; *Euclidean* iff for all $x, y, z \in X$, if $R(x, y)$ and $R(x, z)$ then $R(y, z)$; an *equivalence relation* iff it is reflexive, transitive, and Euclidean. Finally, for any $Y, Z \subseteq X$ we let $R(Y, Z)$ mean that for all $y \in Y$ there is a $z \in Z$ such that $R(y, z)$ and for all $z \in Z$ there is a $y \in Y$ such that $R(y, z)$; this is known as *relational lifting*.

The class of relational models is known as \mathcal{K} . The class of relational models where all accessibility relations are equivalence relations is known as $\mathcal{S5}$, and the class of relational models where all accessibility relations are serial, transitive, and Euclidean is known as $\mathcal{KD45}$.

Let two relational models $M = (S, R, V)$ and $M' = (S', R', V')$ be given. A non-empty relation $\mathfrak{R} \subseteq S \times S'$ is a *bisimulation* if for all $(s, s') \in \mathfrak{R}$ and $a \in A$:

atoms $s \in V(p)$ iff $s' \in V'(p)$ for all $p \in P$;

forth for all $t \in S$, if $R_a(s, t)$, then there is a $t' \in S'$ such that $R'_a(s', t')$ and $(t, t') \in \mathfrak{R}$;

back for all $t' \in S'$, if $R'_a(s', t')$, then there is a $t \in S$ such that $R_a(s, t)$ and $(t, t') \in \mathfrak{R}$.

We write $M \Leftrightarrow M'$ (M and M' are *bisimilar*) iff there is a bisimulation between M and M' , and we write $(M, s) \Leftrightarrow (M', s')$ ((M, s) and (M', s') are *bisimilar*) iff there is a bisimulation between M and M' linking s and s' . Similarly, we write $(M, T) \Leftrightarrow (M', T')$ iff there is a bisimulation between M and M' linking every state in T to a state in T' and linking every state in T' to a state in T .

Using the above-defined notion of relational lifting, if \mathcal{M}_1 and \mathcal{M}_2 are sets of pointed models we say that \mathcal{M}_1 and \mathcal{M}_2 are *bisimilar*, denoted $\mathcal{M}_1 \Leftrightarrow \mathcal{M}_2$, if for every $(M_1, s_1) \in \mathcal{M}_1$ there is an $(M_2, s_2) \in \mathcal{M}_2$ such that $(M_1, s_1) \Leftrightarrow (M_2, s_2)$ and for every $(M_2, s_2) \in \mathcal{M}_2$ there is an $(M_1, s_1) \in \mathcal{M}_1$ such that $(M_1, s_1) \Leftrightarrow (M_2, s_2)$.²

We will now define arrow update models. We can think of them as follows. If you remove the valuation from a relational model you get a *relational frame*. We now decorate

²For the purpose of bisimilarity, we could have treated a multi-pointed model (M, T) as a set of pointed models $\{(M, t) \mid t \in T\}$, so that $(M_1, T_1) \Leftrightarrow (M_2, T_2)$ if and only if for every $t_1 \in T_1$ there is a $t_2 \in T_2$ such that $(M_1, t_1) \Leftrightarrow (M_2, t_2)$, and vice versa. As a union of bisimulations is again a bisimulation, that would have defined the same notion as above.

each arrow (pair in the accessibility relation for an agent) with two formulas in some logical language \mathcal{L} : one for a condition that should hold in the source (state) of the arrow and the other that should hold in the target (state) of the arrow. The result is called an *arrow update model*.

Definition 1 (Arrow update model) Given a logical language \mathcal{L} , an *arrow update model* U is a pair (O, RR) where O is a non-empty domain (set) of *outcomes* (also denoted $\mathcal{D}(U)$) and where RR is an *arrow relation* $RR : A \rightarrow \mathcal{P}((O \times \mathcal{L}) \times (O \times \mathcal{L}))$. \dashv

For each agent a , the arrow relation links (outcome, formula) pairs to each other. We write RR_a for $RR(a)$, and we write $(o, \varphi) \rightarrow_a (o', \varphi')$ for $((o, \varphi), (o', \varphi')) \in RR_a$, or even, if the outcomes are unambiguous, $\varphi \rightarrow_a \varphi'$. Formula φ is the *source condition* and formula φ' is the *target condition* of the a -labelled *arrow* from *source* o to *target* o' . A pointed arrow update model, or *arrow update*, is a pair (U, o) where $o \in O$. Similarly, we define the *multi-pointed arrow update model* (U, Q) , where $Q \subseteq O$, known as well as arrow update. There is no confusion with the arrow updates of AUL [20], as those correspond to singleton pointed arrow update models.

Arrow update models are rather similar to the *action models* by Baltag *et al.* [6]. They are compared in Section 6.

2.2 Syntax

We proceed with the language and semantics of *arbitrary arrow update model logic* (AAUML).

Definition 2 (Syntax) The language \mathcal{L} of AAUML consisting of *formulas* φ is inductively defined as

$$\mathcal{L} \ni \varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box_a\varphi \mid [U, o]\varphi \mid [\uparrow]\varphi$$

where $p \in P$, $a \in A$, and where $U = (O, RR)$ with $o \in O$ is an arrow update model with O finite and with RR_a finite for all $a \in A$, and with source and target conditions that are formulas φ . \dashv

The inductive nature of the definition may be unclear from the construct $[U, o]\varphi$. We should think of $[U, o]\varphi$ as an n -ary operator where not only the formula bound by $[U, o]$ is a formula but also all the source and target conditions in U .³ We read $[U, o]\varphi$ as ‘after executing arrow update (U, o) , φ (holds), and $[\uparrow]\varphi$ as ‘after an arbitrary arrow update, φ (holds)’. Other propositional connectives and dual diamond versions of modalities can be defined as usual by abbreviation: $\Diamond_a\varphi := \neg\Box_a\neg\varphi$, $\langle U, o \rangle\varphi := \neg[U, o]\neg\varphi$, and $\langle \uparrow \rangle\varphi := \neg[\uparrow]\neg\varphi$. Expression $\varphi[\psi/p]$ stands for uniform substitution of all occurrences of p in φ for ψ .

³The BNF informatics-style presentation obscures the inductive nature of the language definition, because the source and target conditions of (U, o) are implicit. The mathematics-style presentation of that clause may be clearer:

Let $\varphi \in \mathcal{L}$, let $U = (O, RR)$ be an arrow update with source and target conditions $\varphi_1, \dots, \varphi_n \in \mathcal{L}$ and such that O and RR_a for all $a \in A$ are finite, and let $o \in O$. Then $[U, o]\varphi \in \mathcal{L}$.

A formula is a *modal formula* if it has shape $\Box_a\varphi$, $\Diamond_a\varphi$, $[\uparrow]\varphi$, $\langle\uparrow\rangle\varphi$, $[U, o]\varphi$, or $\langle U, o\rangle\varphi$. The *modal depth* of a formula $\varphi \in \mathcal{L}$ is defined as: $d(p) = 0$, $d(\neg\varphi) = d(\varphi)$, $d(\varphi \wedge \psi) = \max(d(\varphi), d(\psi))$, $d(\Box_a\varphi) = d(\langle\uparrow\rangle\varphi) = d(\varphi) + 1$, and $d([U, o]\varphi) = d(U) + d(\varphi) + 1$, where $d(U)$ is the maximum modal depth of the source and target conditions occurring in U .

The propositional sublanguage is called \mathcal{L}_{pl} (the *propositional formulas*). Adding the *basic modal* construct $\Box_a\varphi$ to \mathcal{L}_{pl} yields \mathcal{L}_{ml} (the language of basic modal logic, the *basic modal formulas*). Additionally adding the construct $[U, o]\varphi$ yields \mathcal{L}_{auml} (the language of arrow update model logic). In the language \mathcal{L} of AAUML, (modalities for) multi-pointed arrow update models are defined by abbreviation as $[U, Q]\varphi := \bigwedge_{o \in Q} [U, o]\varphi$. From here on we also consider such modalities as logical connectives, such that $[U, Q]\varphi$ is a formula in the logical language.

When doing synthesis, we will put formulas in *disjunctive negation normal form*. This fragment DNNF of \mathcal{L} , that is inspired by the disjunctive normal form of propositional logic and the negation normal form of modal logic, is defined as

$$\begin{aligned} \text{DNNF} \ni \quad & \chi ::= \psi \mid (\chi \vee \chi) \\ & \psi ::= \varphi \mid (\psi \wedge \psi) \\ & \varphi ::= p \mid \neg p \mid \Box_a\chi \mid \Diamond_a\chi \mid [U, o]\chi \mid \langle U, o\rangle\chi \mid [\uparrow]\chi \mid \langle\uparrow\rangle\chi \end{aligned}$$

where the source and target conditions in (U, o) are also formulas χ .

This means that a $\varphi \in \mathcal{L}$ is in disjunctive negation normal form if *every subformula* of φ is a disjunction of conjunctions of formulas that are an atom, or the negation of an atom, or that have one of \Box_a , \Diamond_a , $[U, o]$, $\langle U, o\rangle$, $[\uparrow]$ or $\langle\uparrow\rangle$ as main connective. In particular, this means that formulas have to be in DNNF at every modal depth. So, for example, $p \vee \Box(q \vee (\Diamond p \wedge \neg q))$ is in DNNF, while $p \vee \Box(q \vee \neg(\neg\Diamond p \vee q))$ is not.

2.3 Semantics

We continue with the semantics. The semantics are defined by induction on $\varphi \in \mathcal{L}$, and simultaneously with the execution of arrow update models.

Definition 3 (Semantics) Let a relational model $M = (S, R, V)$, a state $s \in S$, an arrow update model $U = (O, RR)$, and a formula $\varphi \in \mathcal{L}$ be given. The truth (or satisfaction) of φ in (M, s) is defined by induction on φ .

$$\begin{aligned} M, s \models p & \quad \text{iff} \quad s \in V(p) \\ M, s \models \neg\varphi & \quad \text{iff} \quad M, s \not\models \varphi \\ M, s \models \varphi \wedge \psi & \quad \text{iff} \quad M, s \models \varphi \text{ and } M, s \models \psi \\ M, s \models \Box_a\varphi & \quad \text{iff} \quad M, t \models \varphi \text{ for all } (s, t) \in R_a \\ M, s \models [U, o]\varphi & \quad \text{iff} \quad M * U, (s, o) \models \varphi \text{ where } M * U \text{ is defined in } (\#) \\ M, s \models [\uparrow]\varphi & \quad \text{iff} \quad M, s \models [U, o]\varphi \text{ for all } (U, o) \text{ satisfying } (\#\#) \end{aligned}$$

(#): $M * U = (S', R', V')$ is defined as

$$\begin{aligned}
S' &= S \times O \\
&\text{For all } a \in A, \varphi, \varphi' \in \mathcal{L}, s, s' \in S, o, o' \in O : \\
((s, o), (s', o')) \in R'_a &\text{ iff } (s, s') \in R_a, (o, o') \rightarrow_a (o', o'), M, s \models \varphi, \text{ and } M, s' \models \varphi' \\
&\text{For all } p \in P : \\
V'(p) &= V(p) \times O
\end{aligned}$$

(##): (U, o) is an arrow update with all source and target conditions in \mathcal{L}_{ml} .

Formula φ is *valid in* M , notation $M \models \varphi$, iff $M, s \models \varphi$ for all $s \in S$; and φ is *valid* iff for all relational models M we have that $M \models \varphi$. Formulas $\varphi, \psi \in \mathcal{L}$ are *equivalent* iff for all $M = (S, R, V)$ and for all $s \in S$, $M, s \models \varphi$ iff $M, s \models \psi$. The set of validities, also more properly known as the *logic*, is called AAUML. \dashv

Formulas φ and ψ from different languages will also be called equivalent if they satisfy the above condition. The term AAUML will also continue to be informally used for arbitrary arrow update logic. The restriction of arrow formulas to \mathcal{L}_{ml} in the semantics of $[\uparrow]\varphi$ is to avoid circularity of the semantics, as $[\uparrow]\varphi$ could otherwise itself be one of those arrow formulas. However, because we will prove that AAUML is as expressive as basic modal logic, we also have

$$M, s \models [\uparrow]\varphi \quad \text{iff} \quad M, s \models [U, o]\varphi \text{ for all } (U, o)$$

without any restriction on the source and target conditions of U . We will prove this property in Proposition 16, later.

We conclude this subsection by noting two relatively simple properties of AAUML that will be useful in later sections. Firstly, AAUML is invariant under bisimulation, i.e., if $(M, s) \Leftrightarrow (M', s')$ then for all $\varphi \in \mathcal{L}$ we have $M, s \models \varphi$ iff $M', s' \models \varphi$. In [32, Lemma 3] it was shown that AUML is invariant under bisimulation. The proof given in [32] can easily be extended to a proof that AAUML is also invariant under bisimulation.

Secondly, every formula $\varphi \in \mathcal{L}$ is equivalent to a formula φ' that is in DNNF. Proving the existence of such φ' is conceptually simple but rather notationally complex. We therefore provide only an example, and trust that the reader can see that the demonstrated process can be generalized to any $\varphi \in \mathcal{L}$. Suppose that $\varphi = p \wedge \neg(\Box_a \psi_1 \wedge \neg[U, o]\psi_2)$. Our first step is to treat the non-propositional subformulas of φ as atoms, i.e., we treat the formula as $p \wedge \neg(q_1 \wedge \neg q_2)$. This is a formula of propositional logic, so it is equivalent to a formula in disjunctive normal form: $(p \wedge \neg q_1) \vee (p \wedge q_2)$. Then we recall what q_1 and q_2 represent, and obtain $(p \wedge \neg \Box_a \psi_1) \vee (p \wedge [U, o]\psi_2)$. Using the fact that $\neg \Box_a \psi_1$ is equivalent to $\Diamond_a \neg \psi_1$, we find that φ is equivalent to $(p \wedge \Diamond_a \neg \psi_1) \vee (p \wedge [U, o]\psi_2)$. We then repeat this process for $\neg \psi_1$, ψ_2 and every formula χ that occurs as a source or target condition in U . The depth of $\neg \psi_1$, ψ_2 and every such χ is strictly lower than that of φ , so this process eventually terminates, resulting in a formulas ψ'_1, ψ'_2 and χ' that are in DNNF and equivalent to $\neg \psi_1, \psi_2$ and χ , respectively. Let U' be the result of replacing every χ in U by the equivalent χ' . Then the formula $(p \wedge \Diamond_a \psi'_1) \vee (p \wedge [U', o]\psi'_2)$ is in DNNF and equivalent to φ .

2.4 Example

First consider the action of the introductory section of Anne reading a letter containing the truth about p while Bill remains uncertain whether she performs that action. The arrow update model producing the resulting information state is depicted in the upper part of Figure 2. In the figure, an arrow \rightarrow labelled with $\varphi \rightarrow_i \varphi'$ and linking outcomes o, o' stands for the arrow $\varphi \rightarrow_i \varphi'$ between these outcomes, i.e., $((o, \varphi), (o', \varphi')) \in RR_i$; $\varphi \rightarrow_{ij} \varphi'$ stands for $\varphi \rightarrow_i \varphi'$ and $\varphi \rightarrow_j \varphi'$.

In the resulting model Bill considers it possible that Anne knows p , that she knows $\neg p$, and that she still is uncertain about p : $\diamond_b \Box_a p \wedge \diamond_b \Box_a \neg p \wedge \diamond_b \neg(\Box_a p \vee \Box_a \neg p)$.

Next, consider the action of Anne privately learning that p while Bill remains unaware of her doing so. The arrow update model achieving that and the resulting relational model are depicted in the lower part of Figure 2. In the resulting model it is true that, for example, Anne believes p but Bill incorrectly believes that Anne is uncertain about p : $\Box_a p \wedge \Box_b \neg(\Box_a p \vee \Box_a \neg p)$.




Figure 2: Different ways of Anne learning that p

The relation RR_a allows for multiple pairs between the same outcomes. This is necessary. For an example, the singleton arrow update with two reflexive arrows $p \rightarrow_a q, r \rightarrow_a s$ (i.e., $(o, p)RR_a(o, q)$ and $(o, r)RR_a(o, s)$) does not correspond to an arrow update where for any given agent a at most a single arrow links any two outcomes, see [20, 21, 32].

Arrow updates apply to any kind of relational model, and also in particular to relational models wherein all accessibility relations are equivalence relations, the class $\mathcal{S5}$. These relations model *knowledge* of an agent. Relational models wherein all accessibility relations are serial, transitive, and Euclidean, are of the class $\mathcal{KD45}$. These relations model *consistent belief* of an agent. As dynamic epistemic logics typically formalize change of knowledge or change of belief, i.e., *epistemic change*, of particular interest are therefore arrow updates that are $\mathcal{S5}$ -preserving or $\mathcal{KD45}$ -preserving, by which we mean that, given a relational

model in class $\mathcal{S5}$, the update will produce a relational model in class $\mathcal{S5}$, and similarly for $\mathcal{KD45}$.

The examples in this section are indeed typical in that sense. The first example is an $\mathcal{S5}$ -preserving update and the second example is a $\mathcal{KD45}$ -preserving update.

There is more to be learnt from these examples: the first arrow update ‘seems $\mathcal{S5}$ ’ and the second update ‘seems $\mathcal{KD45}$ ’. It is easy to make ‘seem’ precise: consider the following accessibility relation between outcomes induced by an arrow relation:

$$o \rightarrow'_a o' \quad \text{iff} \quad (o, \varphi) \rightarrow_a (o', \varphi') \text{ for some } \varphi, \varphi'.$$

Let an arrow update be *in class* $\mathcal{S5}$ if for all agents a , these induced \rightarrow'_a are equivalence relations; and similarly for $\mathcal{KD45}$. The first arrow update is therefore $\mathcal{S5}$ and the second arrow update is $\mathcal{KD45}$. However, an $\mathcal{S5}$ arrow update of an $\mathcal{S5}$ relational model may not result in an $\mathcal{S5}$ relational model (whereas an $\mathcal{S5}$ action model executed in an $\mathcal{S5}$ relational model will always result in an $\mathcal{S5}$ relational model). This is obvious, as the presence of arrows in the resulting model is also determined by source and target conditions. For example, if in the arrow update of the first example we change the arrow $\top \rightarrow_b \top$ linking \circ to \bullet into $\perp \rightarrow_b \top$, then the resulting model will no longer be reflexive. It is no longer $\mathcal{S5}$. It is not known how to address such issues systematically (see Section 8).

As said, arrow updates are an alternative modelling mechanism to the better known action models. In Section 6, and in particular Subsection 6.5, we compare the two mechanisms in more detail, we will give action models that define the same update as the arrow updates in this section, and we will also present typical applications on which they perform differently.

3 Arrow update synthesis

The goal of synthesis is to find, given a goal formula φ , an arrow update (i.e., pointed arrow update model) (U, o) that makes φ true. There are at least three ways in which we could interpret this goal, however.

Definition 4 (Synthesis)

- The *local synthesis problem* takes as input a pointed model (M, s) and a goal formula φ . The output is an arrow update (U, o) such that $M, s \models \langle U, o \rangle \varphi$, or “NO” if no such arrow update exists.
- The *valid synthesis problem* takes as input a goal formula φ . The output is an arrow update (U, o) such that $\models \langle U, o \rangle \varphi$, or “NO” if no such arrow update exists.
- The *global synthesis problem* takes as input a goal formula φ . The output is an arrow update (U, o) such that for every pointed model (M, s) , if there is some (U', o') such that $M, s \models \langle U', o' \rangle \varphi$, then $M, s \models \langle U, o \rangle \varphi$.

We recall from the introduction that we take the third approach: when we say synthesis we mean global synthesis. An alternative, equivalent characterization of the global synthesis problem is that, for given φ , we want to find (U, o) such that, for all (M, s) , $M, s \models \langle \uparrow \rangle \varphi \leftrightarrow \langle U, o \rangle \varphi$ (see Proposition 16). We further recall that synthesis is impossible for PAL and for AUL, but possible for AML [14]. We now show that synthesis for AUML is also possible. Because our version of synthesis is global, it cannot depend on any specific model. So our synthesis process is purely syntactic.

In our synthesis, we make use of so-called *reduction axioms*. These reduction axioms are a set of validities that, when taken together, show that AAUML has the same expressive power as modal logic.

3.1 Reduction axioms for arrow update models

We start by considering the reduction axioms for the $[U, o]$ operator.

Lemma 5 ([21]) Let (U, o) be an arrow update, $p \in P$, $a \in A$ and $\varphi, \psi \in \mathcal{L}$. Then the following validities hold.

$$\begin{aligned}
& \models [U, o]p \leftrightarrow p \\
& \models [U, o]\neg\varphi \leftrightarrow \neg[U, o]\varphi \\
& \models [U, o](\varphi \wedge \psi) \leftrightarrow ([U, o]\varphi \wedge [U, o]\psi) \\
& \models [U, o]\Box_a\varphi \leftrightarrow \bigwedge_{(o, \psi) \rightarrow_a (o', \psi')} (\psi \rightarrow \Box_a(\psi' \rightarrow [U, o']\varphi))
\end{aligned}$$

Proof The first three validities follow immediately from the semantics of $[U, o]$. The fourth validity also follows from the semantics, in the following way.

$$\begin{aligned}
M, w \models [U, o]\Box_a\varphi & \text{ iff } M * U, (w, o) \models \Box_a\varphi \\
& \text{ iff for all } (w', o') \text{ such that } (w, o)R_a(w', o') : M * U, (w', o') \models \varphi \\
& \text{ iff for all } (o', \psi') \text{ and } w' \text{ such that } (o, \psi) \rightarrow_a (o', \psi') \text{ and } R_a w w' : \\
& \quad \text{if } M, w \models \psi \text{ and } M, w' \models \psi' \text{ then } M * U, (w', o') \models \varphi \\
& \text{ iff for all } (o', \psi') \text{ such that } (o, \psi) \rightarrow_a (o', \psi') : \\
& \quad \text{if } M, w \models \psi \text{ then } M, w \models \Box_a(\psi' \rightarrow [U, o']\varphi) \\
& \text{ iff } M, w \models \bigwedge_{(o, \psi) \rightarrow_a (o', \psi')} (\psi \rightarrow \Box_a(\psi' \rightarrow [U, o']\varphi)) \quad \square
\end{aligned}$$

Note that, in particular, $\models [U, o]\neg\varphi \leftrightarrow \neg[U, o]\varphi$ implies that $[U, o]$ is self-dual: we have $\models [U, o]\varphi \leftrightarrow \langle U, o \rangle \varphi$. This, of course, does not extend to the arbitrary arrow update operator: there are φ for which $\not\models [\uparrow]\varphi \leftrightarrow \langle \uparrow \rangle \varphi$.

The above lemma shows that $[U, o]$ commutes with \neg , distributes over \wedge and, in a somewhat complicated way, commutes with \Box_a . As discussed in [21], this suffices to show that $[U, o]$ can be eliminated from the restriction of the language \mathcal{L} to \mathcal{L}_{auml} .

Corollary 6 For every $\varphi \in \mathcal{L}_{auml}$ there is a formula $\varphi' \in \mathcal{L}_{ml}$ such that $\models \varphi \leftrightarrow \varphi'$. \dashv

3.2 Reduction axioms for the arrow update model quantifier

We can also write similar reduction axioms for $\langle \uparrow \rangle$. In practice, however, it turns out to be slightly more convenient to write them for the dual operator $\langle \uparrow \rangle$. Note that in the lemmas in this subsection we restrict ourselves to the language \mathcal{L}_{auml} , as some of those lemmas use that $\langle \uparrow \rangle$ quantifies over arrow updates with source and target conditions in \mathcal{L}_{ml} , and because we can meet this constraint by applying Corollary 6. Later, in Theorem 15 in the next subsection, we will show that this restriction is unnecessary, and that the lemmas apply to \mathcal{L} as well.

Lemma 7 For every $\varphi \in \mathcal{L}_{auml}$ and every $a \in A$, we have

$$\models \langle \uparrow \rangle \diamond_a \varphi \leftrightarrow \diamond_a \langle \uparrow \rangle \varphi \quad \dashv$$

Proof Let (M, w) be any pointed model, and suppose that $M, w \models \langle \uparrow \rangle \diamond_a \varphi$. Then there is some (U, o) such that $M * U, (w, o) \models \diamond_a \varphi$. So (w, o) has an a -successor (w', o') such that $M * U, (w', o') \models \varphi$.

This implies that $M, w' \models \langle U, o' \rangle \varphi$ and therefore $M, w' \models \langle \uparrow \rangle \varphi$. Since w' is an a -successor of w , we obtain $M, w \models \diamond_a \langle \uparrow \rangle \varphi$.

Now, suppose that $M, w \models \diamond_a \langle \uparrow \rangle \varphi$. Then there is an a -successor w' of w such that $M, w' \models \langle \uparrow \rangle \varphi$. As witness for this $\langle \uparrow \rangle$ statement there must be some U', o' such that $M, w' \models \langle U', o' \rangle \varphi$.

Let (U, o) be the arrow update obtained by adding one extra world o to U' , and a transitions $(o, \top) \rightarrow_a (o', \top)$. Note that $(M * U', (w', o'))$ is bisimilar to $(M * U, (w', o'))$, and therefore $M * U, (w', o') \models \varphi$. Finally, note that (w', o') is an a -successor of (w, o) , so we have $M * U, (w, o) \models \diamond_a \varphi$ and therefore $M, w \models \langle \uparrow \rangle \diamond_a \varphi$. \square

Note that the proof is constructive. That is, if we find (U', o') such that $M, w \models \diamond_a \langle U', o' \rangle \varphi$ then not only do we know that $M, w \models \langle \uparrow \rangle \diamond_a \varphi$, we can also find a specific (U, o) such that $M, w \models \langle U, o \rangle \diamond_a \varphi$.

Next, we consider a slightly stronger lemma.

Lemma 8 For every $\varphi_1, \dots, \varphi_n \in \mathcal{L}_{auml}$ and every $a \in A$ we have

$$\models \langle \uparrow \rangle \bigwedge_{1 \leq i \leq n} \diamond_a \varphi_i \leftrightarrow \bigwedge_{1 \leq i \leq n} \diamond_a \langle \uparrow \rangle \varphi_i$$

Proof The left-to-right direction is obvious, so we show only the right-to-left direction. So suppose that $M, w \models \bigwedge \diamond_a \langle \uparrow \rangle \varphi_i$. Then there are a -successors w_1, \dots, w_n of w and pointed arrow update models $(U_1, o_1), \dots, (U_n, o_n)$ such that $M, w_i \models \langle U_i, o_i \rangle \varphi_i$ for all i .

Now, let (U, o) be the arrow update obtained by taking the disjoint union of all U_i and adding one extra outcome o , and adding arrows $(o, \top) \rightarrow_a (o_i, \top)$ for every o_i .

For every i , $(M * U_i, (w_i, o_i))$ is bisimilar to $(M * U, (w_i, o_i))$, so we have $M * U, (w_i, o_i) \models \varphi_i$. Finally, (w_i, o_i) is an a -successor of (w, o) for every i . As such, we have $M, w \models \langle U, o \rangle \bigwedge \diamond_a \varphi_i$ and therefore, as all the source and target conditions of U are in \mathcal{L}_{ml} , $M, w \models \langle \uparrow \rangle \bigwedge \diamond_a \varphi_i$. \square

Again, the proof is constructive, so given (U_i, o_i) for all i , we can find the model (U, o) . Note also that the φ_i need not be consistent with each other.

Some reflection may be in order as to why Lemma 8 holds. Suppose that $M, w \models \bigwedge \diamond_a \langle \uparrow \rangle \varphi_i$. So for every i , there is some world w_i that a considers possible as well as some event U_i and outcome o_i such that, if (U_i, o_i) were to happen in w_i , then φ_i would become true.

Now let us look at the arrow update (U, o) that we constructed. Effectively, this arrow update represents us telling a that “we are performing one of the actions U_i, o_i , but we are not telling you which one.” Now, for every i agent a considers it possible that w_i is the actual world, and that (U_i, o_i) is the event that happened. As such, after we execute our event we are in a situation where every φ_i is held possible by a .

So far, we have only considered diamonds. Now, let us add a box modality.

Lemma 9 For every $\varphi_1, \dots, \varphi_n, \psi \in \mathcal{L}_{auml}$ and every $a \in A$, we have

$$\models \langle \uparrow \rangle \left(\bigwedge_{1 \leq i \leq n} \diamond_a \varphi_i \wedge \Box_a \psi \right) \leftrightarrow \bigwedge_{1 \leq i \leq n} \diamond_a \langle \uparrow \rangle (\varphi_i \wedge \psi)$$

Proof The left-to-right direction is fairly obvious. Suppose that $M, w \models \langle \uparrow \rangle (\bigwedge_{1 \leq i \leq n} \diamond_a \varphi_i \wedge \Box_a \psi)$. Then there is a (U, o) such that $M, w \models \langle U, o \rangle (\bigwedge_{1 \leq i \leq n} \diamond_a \varphi_i \wedge \Box_a \psi)$. Therefore, $M * U, (w, o) \models \Box_a \psi$ and for each $1 \leq i \leq n$, $M * U, (w, o) \models \diamond_a \varphi_i$. Let (w_i, o') be such that $(w, o) R_a (w_i, o')$ and $M * U, (w_i, o') \models \varphi_i$. From $M * U, (w, o) \models \Box_a \psi$ and $(w, o) R_a (w_i, o')$ also follows that $M * U, (w_i, o') \models \psi$. Combining both we have $M * U, (w_i, o') \models \varphi_i \wedge \psi$. Therefore, $M, w_i \models \langle U, o' \rangle (\varphi_i \wedge \psi)$, from which it follows that $M, w_i \models \langle \uparrow \rangle (\varphi_i \wedge \psi)$. From $(w, o) R_a (w_i, o')$ it follows by definition that $w R_a w_i$. From $M, w_i \models \langle \uparrow \rangle (\varphi_i \wedge \psi)$ and $w R_a w_i$ we get the required $M, w \models \diamond_a \langle \uparrow \rangle (\varphi_i \wedge \psi)$. As i was arbitrary, $M, w \models \bigwedge_{1 \leq i \leq n} \diamond_a \langle \uparrow \rangle (\varphi_i \wedge \psi)$.

We now show the right-to-left direction. So suppose that $M, w \models \bigwedge_{1 \leq i \leq n} \diamond_a \langle \uparrow \rangle (\varphi_i \wedge \psi)$. Then for every $1 \leq i \leq n$, there are an a -successor w_i of w and (U_i, o_i) such that $M, w_i \models \langle U_i, o_i \rangle (\varphi_i \wedge \psi)$.

Let (U, o) be the model obtained by taking the disjoint union of all U_i , and adding a single outcome o with arrows $(o, \top) \rightarrow_a (o_i, \langle U_i, o_i \rangle \psi)$ for every i .

Consider $(M * U, (w, o))$. By assumption, $M, w_i \models \langle U_i, o_i \rangle (\varphi_i \wedge \psi)$, so $M, w_i \models \langle U_i, o_i \rangle \psi$. From that and the fact that model U contains arrow $(o, \top) \rightarrow_a (o_i, \langle U_i, o_i \rangle \psi)$ it follows that (w_i, o_i) is an a -successor of (w, o) in $(M * U)$. Furthermore, also from $M, w_i \models \langle U_i, o_i \rangle (\varphi_i \wedge \psi)$ it follows that $M, w_i \models \langle U_i, o_i \rangle \varphi_i$, i.e., $M * U_i, (w_i, o_i) \models \varphi_i$, and as $(M * U, (w_i, o_i))$ is bisimilar to $(M * U_i, (w_i, o_i))$ it follows that $M * U, (w_i, o_i) \models \varphi_i$ (the bisimulation is the identity relation on the restriction of the domain of $M * U$ to the domain of $M * U_i$). With $(w, o) R_a (w_i, o_i)$ we thus get $M * U, (w, o) \models \diamond_a \varphi_i$, and as i was arbitrary, $M * U, (w, o) \models \bigwedge_{1 \leq i \leq n} \diamond_a \varphi_i$.

Additionally, note that every outgoing a -arrow in (U, o) has target condition $\langle U_i, o_i \rangle \psi$ for some i . Again observing that $(M * U, (w_i, o_i))$ is bisimilar to $(M * U_i, (w_i, o_i))$, it follows that for every (w_i, o_i) that is an a -successor of (w, o) , we have $M * U, (w_i, o_i) \models \psi$. It follows that $M * U, (w, o) \models \Box_a \psi$.

Taken together, the above shows that $M * U, (w, o) \models \bigwedge_{1 \leq i \leq n} \diamond_a \varphi_i \wedge \square_a \psi$ and thus that $M, w \models \langle U, o \rangle (\bigwedge_{1 \leq i \leq n} \diamond_a \varphi_i \wedge \square_a \psi)$. Furthermore, every formula χ occurring in U is either a basic modal formula, or of the form $\langle U_i, o_i \rangle \psi$. Because $\psi \in \mathcal{L}_{auml}$, we have $\langle U_i, o_i \rangle \psi \in \mathcal{L}_{auml}$ and therefore, by Corollary 6, there is a formula $\chi_i \in \mathcal{L}_{ml}$ that is equivalent to $\langle U_i, o_i \rangle \psi$. Let U' be the arrow update model obtained from U by replacing $\langle U_i, o_i \rangle \psi$ by χ_i , for every i . Then we also have $M, w \models \langle U', o \rangle (\bigwedge_{1 \leq i \leq n} \diamond_a \varphi_i \wedge \square_a \psi)$. Because U' only contains formulas of modal logic, it follows that $M, w \models \langle \uparrow \rangle (\bigwedge_{1 \leq i \leq n} \diamond_a \varphi_i \wedge \square_a \psi)$, as was to be shown. \square

Once again, the proof is constructive. Note that on the right-hand side we have eliminated the \square_a connective. This is a consequence of the fact that as the designer of the arrow update model U , we have the freedom to inform a that certain worlds, which she might previously have considered possible, are not in fact the actual world. This results in the removal of the a -arrows to these worlds. So if we want to make $\square_a \psi$ true after the application of U , then we can simply have a eliminate all successors where ψ would otherwise become false. In the construction used in the lemma, we do this using the target condition $\langle U_i, o_i \rangle \psi$.

In the preceding three lemmas, we only considered \diamond_a and \square_a operators for one single agent a . However, when constructing (U, o) we can place arrows for different agents independently, so the same construction works for multiple agents at the same time. This yields the following lemma.

Lemma 10 For every $a \in A$, let $\Phi_a \subseteq \mathcal{L}_{auml}$ be a finite set of formulas, and let $\psi_a \in \mathcal{L}_{auml}$ be a formula. Then

$$\models \langle \uparrow \rangle \bigwedge_{a \in A} \left(\bigwedge_{\varphi_a \in \Phi_a} \diamond_a \varphi_a \wedge \square_a \psi_a \right) \leftrightarrow \bigwedge_{a \in A} \bigwedge_{\varphi_a \in \Phi_a} \diamond_a \langle \uparrow \rangle (\varphi_a \wedge \psi_a)$$

Proof We proceed as in the previous Lemma 9, but simultaneously for all $a \in A$.

Again, the left-to-right direction is fairly obvious, and completely analogous to the treatment in Lemma 9. Suppose that $M, w \models \langle \uparrow \rangle \bigwedge_{a \in A} (\bigwedge_{\varphi_a \in \Phi_a} \diamond_a \varphi_a \wedge \square_a \psi_a)$. Then there is a (U, o) such that $M, w \models \langle U, o \rangle (\bigwedge_{a \in A} (\bigwedge_{\varphi_a \in \Phi_a} \diamond_a \varphi_a \wedge \square_a \psi_a))$. Therefore, for all $a \in A$ and for all $\varphi_a \in \Phi_a$, $M * U, (w, o) \models \square_a \psi$ and $M * U, (w, o) \models \diamond_a \varphi_a$. As before we now obtain $M, w \models \diamond_a \langle \uparrow \rangle (\varphi_a \wedge \psi)$, for all $a \in A$ and for all $\varphi_a \in \Phi_a$, and thus $M, w \models \bigwedge_{a \in A} \bigwedge_{\varphi_a \in \Phi_a} \diamond_a \langle \uparrow \rangle (\varphi_a \wedge \psi)$.

From right to left, suppose that $M, w \models \bigwedge_{a \in A} \bigwedge_{\varphi_a \in \Phi_a} \diamond_a \langle \uparrow \rangle (\varphi_a \wedge \psi_a)$. Let $\Phi_a = \{\varphi_a^1, \dots, \varphi_a^{|\Phi_a|}\}$ (Φ_a may be empty). Then for every $a \in A$ and for every $1 \leq i \leq |\Phi_a|$, there are (U_a^i, o_a^i) and an a -successor w_a^i of w such that $M, w_a^i \models \langle U_a^i, o_a^i \rangle (\varphi_a^i \wedge \psi_a)$. Similarly to the previous lemma, we let (U, o) be the model obtained by taking the disjoint union of all U_a^i (i.e., for all $a \in A$ and for all $1 \leq i \leq |\Phi_a|$), but still only adding a *single* outcome o with arrows $(o, \top) \rightarrow_a (o_i, \langle U_i, o_i \rangle \psi)$ for every $a \in A$ and $1 \leq i \leq |\Phi_a|$. We then proceed as before. \square

Lemma 10 is the most important reduction axiom for AAUML. However, not every formula is of a form such that the lemma can be applied. We therefore need two validities that allow us to put formulas in the correct form.

Lemma 11 For every $\varphi_1, \varphi_2 \in \mathcal{L}_{auml}$ and every $\varphi_0 \in \mathcal{L}_{pl}$, we have

$$\models \langle \uparrow \rangle (\varphi_1 \vee \varphi_2) \leftrightarrow (\langle \uparrow \rangle \varphi_1 \vee \langle \uparrow \rangle \varphi_2)$$

and

$$\models \langle \uparrow \rangle (\varphi_0 \wedge \varphi_1) \leftrightarrow (\varphi_0 \wedge \langle \uparrow \rangle \varphi_1).$$

Proof The proof of this lemma is straightforward.

By the semantic definition, from $M, w \models \langle \uparrow \rangle (\varphi_1 \vee \varphi_2)$ it follows that $M, w \models \langle \uparrow \rangle \varphi_1$ or that $M, w \models \langle \uparrow \rangle \varphi_2$, and therefore that $M, w \models \langle \uparrow \rangle \varphi_1 \vee \langle \uparrow \rangle \varphi_2$. In the other direction, from $M, w \models \langle \uparrow \rangle \varphi_1 \vee \langle \uparrow \rangle \varphi_2$ it follows that $M, w \models \langle \uparrow \rangle \varphi_1$ or that $M, w \models \langle \uparrow \rangle \varphi_2$, and therefore, by weakening either formula to the disjunction $\varphi_1 \vee \varphi_2$, that $M, w \models \langle \uparrow \rangle (\varphi_1 \vee \varphi_2)$.

For the second validity, we first observe that $\langle \uparrow \rangle \varphi_0 \leftrightarrow \varphi_0$ (*) is a validity of AAUML (all our logics are logics of informational change, not of factual change). From $M, w \models \langle \uparrow \rangle (\varphi_0 \wedge \varphi_1)$ now follows $M, w \models \langle \uparrow \rangle \varphi_0$ and $M, w \models \langle \uparrow \rangle \varphi_1$, and thus, using (*), that $M, w \models \varphi_0$ and $M, w \models \langle \uparrow \rangle \varphi_1$, and so $M, w \models \varphi_0 \wedge \langle \uparrow \rangle \varphi_1$. For the other direction we observe that, on the assumption of $M, w \models \varphi_0 \wedge \langle \uparrow \rangle \varphi_1$, any arrow update executed in (M, w) to make φ_1 true will preserve, by (*), the truth of φ_0 , so that $M, w \models \langle \uparrow \rangle (\varphi_0 \wedge \varphi_1)$. \square

It is important and non-trivial to note that the disjunction case can be made constructive. Suppose that we have already synthesized (U_1, o_1) and (U_2, o_2) such that $\models \langle \uparrow \rangle \varphi_1 \leftrightarrow \langle U_1, o_1 \rangle \varphi_1$ and $\models \langle \uparrow \rangle \varphi_2 \leftrightarrow \langle U_2, o_2 \rangle \varphi_2$. So we have two pointed arrow update models that make φ_1 and φ_2 true whenever possible. This does not, however, immediately give us a *single*-pointed arrow update model (U, o) that guarantees $\varphi = \varphi_1 \vee \varphi_2$ whenever possible.⁴ In order to find this (U, o) , we have to combine (U_1, o_1) and (U_2, o_2) . We do this in the following way.

First, we take the set of outcomes of U to be the disjoint union of the sets of outcomes of U_1 and U_2 , plus one extra outcome o . Then, we add arrows as follows to U .

For every $(o_1, \psi) \rightarrow_a (o', \psi')$ of U_1 , add an arrow $(o, \psi \wedge \langle U_1, o_1 \rangle \varphi_1) \rightarrow_a (o', \psi')$.

For every $(o_2, \psi) \rightarrow_a (o', \psi')$ of U_2 , add an arrow $(o, \psi \wedge \neg \langle U_1, o_1 \rangle \varphi_1) \rightarrow_a (o', \psi')$.

When executed in a $\langle \uparrow \rangle \varphi_1$ world, this arrow update (U, o) will act as (U_1, o_1) , since every such world satisfies $\langle U_1, o_1 \rangle \varphi_1$ and we added all arrows from o_1 with an extra $\langle U_1, o_1 \rangle \varphi_1$ precondition. When executed in any $\neg \langle \uparrow \rangle \varphi_1$ world, (U, o) acts as (U_2, o_2) . As long as either $\langle \uparrow \rangle \varphi_1$ or $\langle \uparrow \rangle \varphi_2$ holds, we therefore have $\langle U, o \rangle (\varphi_1 \vee \varphi_2)$.

More formally, we have the following lemma.

Lemma 12 Let $\varphi_1, \varphi_2 \in \mathcal{L}_{auml}$, and for $i = 1, 2$ let (U_i, o_i) be such that $\models \langle \uparrow \rangle \varphi_i \leftrightarrow \langle U_i, o_i \rangle \varphi_i$, where $U_i = (O_i, RR_i)$. Furthermore, let $U = (O, RR)$ be given as follows:

⁴An alternative technique to synthesize an arrow update for the disjunction is to take the double-pointed direct sum of (U_1, o_1) and (U_2, o_2) ; its points are $\{o_1, o_2\}$. This method is followed in [14].

- $O = \{o\} \uplus O_1 \uplus O_2$,
- RR contains exactly the following arrows:
 1. $(o', \psi') \rightarrow_a (o'', \psi'') \in RR_i$, for $i = 1, 2$,
 2. $(o, \psi \wedge \langle U_1, o_1 \rangle \varphi_1) \rightarrow_a (o', \psi')$ where $(o_1, \psi) \rightarrow_a (o', \psi') \in RR_1$,
 3. $(o, \psi \wedge \neg \langle U_1, o_1 \rangle \varphi_1) \rightarrow_a (o', \psi')$ where $(o_2, \psi) \rightarrow_a (o', \psi') \in RR_2$.

Then $\models \langle \uparrow \rangle (\varphi_1 \vee \varphi_2) \leftrightarrow \langle U, o \rangle (\varphi_1 \vee \varphi_2)$. ⊖

Proof The right-to-left direction is obvious, because, just as in the proof of Lemma 9, source conditions $\psi \wedge \langle U_1, o_1 \rangle \varphi_1$ and $\psi \wedge \neg \langle U_1, o_1 \rangle \varphi_1$ can with Corollary 6 be assumed to be equivalent to modal logical formulas.

We now show the left-to-right direction. Suppose therefore that $M, w \models \langle \uparrow \rangle (\varphi_1 \vee \varphi_2)$. Then, by Lemma 11, we have $M, w \models \langle \uparrow \rangle \varphi_1 \vee \langle \uparrow \rangle \varphi_2$. We continue by a case distinction.

First, suppose that $M, w \models \langle \uparrow \rangle \varphi_1$. Then none of the $(o, \psi \wedge \neg \langle U_1, o_1 \rangle \varphi_1) \rightarrow_a (o', \psi')$ arrows are applicable in w . The $(o, \psi \wedge \langle U_1, o_1 \rangle \varphi_1) \rightarrow_a (o', \psi')$ arrows, on the other hand, are applicable if and only if $(o_1, \psi) \rightarrow_a (o', \psi')$ applies. Therefore, for every agent a , $(w, o)R_a(w', o')$ iff $(wR_a w'$ and $oR_a o')$ iff $(wR_a w'$ and $o_1R_a o')$ iff $(w, o_1)R_a(w', o')$: there is an a -arrow from (w, o) to (w', o') in $M * U$ if and only if there is an a -arrow from (w, o_1) to (w', o') in $M * U_1$. Furthermore, because U contains a copy of U_1 , we have that for every $o', o'' \in O_1$ and every $w', w'' \in W$ there is an a -arrow from (w', o') to (w'', o'') in $M * U$ if and only if there is such an arrow in $M * U_1$. It follows that the relation $\mathfrak{R}_1 = \{((w', o'), (w', o')) \mid w' \in W, o' \in O_1\} \cup \{((w, o), (w, o_1))\}$ is a bisimulation between $M * U, (w, o)$ and $M * U_1, (w, o_1)$. (It also obviously satisfies **atoms**, as this depends on correspondence of the first argument in each (state, outcome) pair.) By assumption we have $M, w \models \langle \uparrow \rangle \varphi_1$ and therefore $M, w \models \langle U_1, o_1 \rangle \varphi_1$, which implies that $M * U_1, (w, o_1) \models \varphi_1$. Because AAUML is invariant under bisimulation it then follows that $M * U, (w, o) \models \varphi_1$ and therefore that $M, w \models \langle U, o \rangle \varphi_1$, so that we also have $M, w \models \langle U, o \rangle (\varphi_1 \vee \varphi_2)$.

Suppose then that $M, w \not\models \langle \uparrow \rangle \varphi_1$. Then we must have $M, w \models \langle \uparrow \rangle \varphi_2$. In this case, the $(o, \psi \wedge \langle U_1, o_1 \rangle \varphi_1) \rightarrow_a (o', \psi')$ arrows are inapplicable while the $(o, \psi \wedge \neg \langle U_1, o_1 \rangle \varphi_1) \rightarrow_a (o', \psi')$ arrows apply if and only if $(o_2, \psi) \rightarrow_a (o', \psi')$ does. Similarly to the previous case, the relation $\mathfrak{R}_2 = \{((w', o'), (w', o')) \mid w' \in W, o' \in O_2\} \cup \{((w, o), (w, o_2))\}$ is therefore a bisimulation between $M * U, (w, o)$ and $M * U_2, (w, o_2)$. From the assumption that $M, w \models \langle \uparrow \rangle \varphi_2$ it follows that $M, w \models \langle U_2, o_2 \rangle \varphi_2$ and therefore that $M * U_2, (w, o_2) \models \varphi_2$. Because AAUML is invariant under bisimulation we then obtain $M * U, (w, o) \models \varphi_2$ and therefore $M, w \models \langle U, o \rangle \varphi_2$ and also $M, w \models \langle U, o \rangle (\varphi_1 \vee \varphi_2)$.

In either case, we have $M, w \models \langle U, o \rangle (\varphi_1 \vee \varphi_2)$, as was to be shown. □

3.3 Reduction

Using the fact that $[U, o]$ commutes with \neg , distributes over \wedge and, in a convoluted but not very complicated way, commutes with \Box_a , Corollary 6 showed that every \mathcal{L}_{auml} formula is

equivalent to a \mathcal{L}_{ml} formula. In order to be able to finally perform synthesis, it remains to show that every \mathcal{L} formula is also equivalent to a \mathcal{L}_{ml} formula. This is what we will show in this section.

In Section 3.2 we showed that $\langle \uparrow \rangle$ commutes, in a very complicated way, with boolean combinations of basic modal formulas. Also, like $[U, o]$, a $\langle \uparrow \rangle$ operator disappears once it encounters a propositional atom. From this we can show that, if φ is a formula of basic modal logic, then we can transform $\langle \uparrow \rangle \varphi$ into an equivalent formula φ' of basic modal logic.

The proof that $\langle \uparrow \rangle \varphi$ can be transformed into φ' is by induction on the order \succ given by: $\varphi_1 \succ \varphi_2$ if and only if (i) φ_2 is a strict subformula of φ_1 or (ii) φ_2 has a (strictly) lower modal depth than φ_1 . Let us first show that this order is well-founded.

Lemma 13 The relation \succ is well-founded. ⊣

Proof This follows from the fact that both the $>$ relation on the natural numbers and the “strict subformula of” relation are well-founded.

Suppose towards a contradiction that there is an infinitely descending chain $\varphi_0 \succ \varphi_1 \succ \dots$ of formulas. So for every $i \in \mathbb{N}$ we have either (i) φ_{i+1} is a strict subformula of φ_i or (ii) $d(\varphi_i) > d(\varphi_{i+1})$. The modal depth of a formula is at least as large as that of its subformulas, so in either case we have $d(\varphi_i) \geq d(\varphi_{i+1})$. Because the depth of a formula is a natural number and $>$ is well-founded on \mathbb{N} , there can be only finitely many i such that $d(\varphi_i) > d(\varphi_{i+1})$. So there is some $k \in \mathbb{N}$ such that for all $m \geq k$ we have $d(\varphi_m) = d(\varphi_k)$.

It follows that for every $m \geq k$, the formula φ_{m+1} must be a strict subformula of φ_m . This is impossible, however, since the “strict subformula of” relation is well-founded. We have arrived at a contradiction, so \succ does not have an infinitely descending sequence and is therefore well-founded. □

Now, we can prove the reduction from $\langle \uparrow \rangle \varphi$ to φ' .

Lemma 14 For every $\varphi \in \mathcal{L}_{ml}$, there is a formula $\varphi' \in \mathcal{L}_{ml}$ such that $\models \langle \uparrow \rangle \varphi \leftrightarrow \varphi'$. ⊣

Proof Every formula is equivalent to a formula in DNNF (See Sections 2.2 and 2.3), so we can assume without loss of generality that φ is in DNNF. We can also assume without loss of generality that every conjunction in φ contains exactly one conjunct of the form $\Box_a \chi$ for every agent a . We now proceed by induction on the order \succ , which, by the previous lemma, is well-founded. Note that the minimal elements with respect to \succ are the formulas that are of depth 0 and have no strict subformulas, i.e., the atoms.

Suppose therefore, as a base case in our induction, that φ is an atom. Then $\models \langle \uparrow \rangle \varphi \leftrightarrow \varphi$, so the lemma holds with $\varphi' = \varphi$. Assume now as induction hypothesis that the lemma holds for every φ' such that $\varphi \succ \varphi'$. We then continue by case distinction.

First, suppose that $\varphi = \varphi_1 \vee \varphi_2$. By Lemma 11 we have $\models \langle \uparrow \rangle (\varphi_1 \vee \varphi_2) \leftrightarrow (\langle \uparrow \rangle \varphi_1 \vee \langle \uparrow \rangle \varphi_2)$. Furthermore, $\varphi \succ \varphi_1$ and $\varphi \succ \varphi_2$, so by the induction hypothesis there are $\varphi'_1, \varphi'_2 \in \mathcal{L}_{ml}$ such that $\models \langle \uparrow \rangle \varphi_1 \leftrightarrow \varphi'_1$ and $\models \langle \uparrow \rangle \varphi_2 \leftrightarrow \varphi'_2$. It follows that $\models \langle \uparrow \rangle \varphi \leftrightarrow (\varphi'_1 \vee \varphi'_2)$.

Secondly, suppose that $\varphi = \varphi_0 \wedge \varphi_1$, where $\varphi_0 \in \mathcal{L}_{pl}$. Then, by Lemma 11, we have $\models \langle \uparrow \rangle (\varphi_0 \wedge \varphi_1) \leftrightarrow (\varphi_0 \wedge \langle \uparrow \rangle \varphi_1)$. Furthermore, $\varphi \succ \varphi_1$, so by the induction hypothesis there is a φ'_1 such that $\models \langle \uparrow \rangle \varphi_1 \leftrightarrow \varphi'_1$. It follows that $\models \langle \uparrow \rangle \varphi \leftrightarrow (\varphi_0 \wedge \varphi'_1)$.

Finally, suppose that φ is a conjunction without a propositional conjunct. Then for every $a \in A$ there are a finite set $\Phi_a \subseteq \mathcal{L}_{ml}$ and a formula $\psi_a \in \mathcal{L}_{ml}$ such that $\varphi = \bigwedge_{a \in A} (\bigwedge_{\varphi_a \in \Phi_a} \diamond_a \varphi_a \wedge \square_a \psi_a)$. By Lemma 10, we have $\models \langle \uparrow \rangle \varphi \leftrightarrow \bigwedge_{a \in A} \bigwedge_{\varphi_a \in \Phi_a} \diamond_a \langle \uparrow \rangle (\varphi_a \wedge \psi_a)$. For every $a \in A$ and $\varphi_a \in \Phi_a$, the modal depth of $\varphi_a \wedge \psi_a$ is strictly lower than that of φ , so $\varphi \succ \varphi_a \wedge \psi_a$. Therefore, by the induction hypothesis, there is a formula $\varphi'_a \in \mathcal{L}_{ml}$ such that $\models \langle \uparrow \rangle (\varphi_a \wedge \psi_a) \leftrightarrow \varphi'_a$. Let Φ'_a be the set of such φ'_a . It follows that $\models \langle \uparrow \rangle \varphi \leftrightarrow \bigwedge_{a \in A} \bigwedge_{\varphi'_a \in \Phi'_a} \diamond_a \varphi'_a$.

This completes the induction step and thereby the proof. \square

We now have the following theorem.

Theorem 15 For every $\varphi \in \mathcal{L}$ there is a formula $\varphi' \in \mathcal{L}_{ml}$ such that $\models \varphi \leftrightarrow \varphi'$. \dashv

Proof We now use Lemma 14 (in dual form) that

for every $\varphi \in \mathcal{L}_{ml}$ there is a $\varphi' \in \mathcal{L}_{ml}$ such that $\models [\uparrow] \varphi \leftrightarrow \varphi'$,

and we also use the consequence of Lemma 14 that

for every $\varphi \in \mathcal{L}_{ml}$ and arrow update (U, o) with source and target conditions in \mathcal{L}_{ml} there is a $\varphi' \in \mathcal{L}_{ml}$ such that $\models [U, o] \varphi \leftrightarrow \varphi'$,

in order to successively eliminate the innermost $[U, o]$ or $[\uparrow]$ operators of any given formula of AAUML. We can thus transform this formula into an equivalent formula of modal logic. \square

It follows that in Lemmas 7–12 the restriction to the sublanguage \mathcal{L}_{auml} is unnecessary; the lemmas also hold for the full language \mathcal{L} . In this form we will also later use these validities as axioms in the axiomatization. For the next section, it is important to keep in mind that the reduction axioms not only guarantee the existence of φ' , but also enable us to find it.

Theorem 15 also allows us to prove a claim that we made in Section 2.3. There, we defined $M, s \models [\uparrow] \varphi$ by

$M, s \models [\uparrow] \varphi$ iff $(M, s \models [U, o] \varphi$ for every arrow update (U, o) that has source and target conditions only in \mathcal{L}_{ml}).

Now that we have shown that every formula of \mathcal{L} is equivalent to a formula of \mathcal{L}_{ml} , it follows immediately that the requirement of the source and target conditions being in \mathcal{L}_{ml} is unnecessary.

Proposition 16 For every $\varphi \in \mathcal{L}$ and every pointed relational model (M, s) , we have

$M, s \models [\uparrow] \varphi$ iff $M, s \models [U, o] \varphi$ for every arrow update (U, o) . \dashv

Procedure $Synth(\varphi)$

Input: $\varphi \in \mathcal{L}$.

Output: (U_φ, o_φ) such that $\models \langle \uparrow \rangle \varphi \leftrightarrow \langle U_\varphi, o_\varphi \rangle \varphi$.

1. If $\varphi \notin \mathcal{L}_{ml}$, then use the reduction axioms to find a formula $\varphi_{modal} \in \mathcal{L}_{ml}$ such that $\models \varphi \leftrightarrow \varphi_{modal}$, and return $Synth(\varphi_{modal})$. Otherwise, continue.
 2. If φ is not in disjunction normal form, compute the DNNF φ_{DNNF} of φ and return $Synth(\varphi_{DNNF})$. Otherwise, continue.
 3. If $\varphi = \varphi_1 \vee \varphi_2$, then compute $Synth(\varphi_1)$ and $Synth(\varphi_2)$, and combine the two arrow update models as in Lemma 12. Return the combined arrow update model.
 4. If φ is not a disjunction, then since it is in DNNF it must be a conjunction, where each conjunct is (i) a literal, (ii) of the form $\diamond_a \psi$, or (iii) of the form $\Box_a \chi$. Assume w.l.o.g. that for every a there is exactly one conjunct $\Box_a \chi_a$. For every $\diamond_a \psi$, compute $Synth(\psi \wedge \chi_a)$. Then use Lemma 9 to combine the arrow update models, and return the result. If there are no \diamond_a operators for any agent a , return the trivial arrow update model with one outcome and no arrows.
-
-

Table 1: Synthesis procedure

3.4 Synthesis

Recall that our goal, when performing synthesis, is to find, for given $\varphi \in \mathcal{L}$, an arrow update (U_φ, o_φ) such that $\models \langle \uparrow \rangle \varphi \leftrightarrow \langle U_\varphi, o_\varphi \rangle \varphi$. Using Theorem 15, we can transform φ into an equivalent formula of modal logic. Then, using the procedure outlined in Section 3.2, we can find (U_φ, o_φ) . The procedure is found in detail in Table 1.

Our arrow update synthesis was motivated by Hales' already mentioned action model synthesis published in [14]. See also the next Section 4. Subsequently Hales *et al.* investigated these matters in [10] and in (his Ph.D. thesis) [15]. Alternatively, action model synthesis by way of a dedicated action language was employed by Aucher in [3].

3.5 Example

In order to gain better understanding of $Synth(\varphi)$, let us consider an example. Suppose $\varphi = \diamond_a \Box_b p \wedge \diamond_b (\diamond_a q \vee \diamond_a r) \wedge \Box_b p$. We want to perform synthesis for this φ .

Goal: find $Synth(\diamond_a \Box_b p \wedge \diamond_b (\diamond_a q \vee \diamond_a r) \wedge \Box_b p)$.

Because $\varphi \in \mathcal{L}_{ml}$, φ is in DNNF and φ is not a disjunction, we continue past steps 1, 2 and 3. In step 4, it is assumed that for every agent there is exactly one \Box conjunct. This means we need to add a trivial conjunct $\Box_a \top$.

Of the conjuncts of φ , two have \diamond as primary operator. So we need to perform synthesis for two more formulas; because of $\diamond_a \Box_b p$ and $\Box_a \top$ we need to find $Synth(\Box_b p \wedge \top)$, and

because of $\diamond_b(\diamond_a q \vee \diamond_a r)$ and $\square_b p$ we need to find $Synth((\diamond_a q \vee \diamond_a r) \wedge p)$.

Subgoal 1: find $Synth(\square_b p \wedge \top)$.

Since we are doing synthesis for a conjunction, we continue in steps 1, 2 and 3. Because there are no \diamond operators in $\square_b p \wedge \top$, we return the trivial arrow update model in step 4.

Subgoal 2: find $Synth((\diamond_a q \vee \diamond_a r) \wedge p)$. In step 2, we need to put the formula in DNNF. We therefore continue with $(\diamond_a q \wedge p) \vee (\diamond_a r \wedge p)$. In step 3 we are then instructed to perform synthesis for the two disjuncts.

Sub-subgoal 2.1: find $Synth(\diamond_a q \wedge p)$.

We continue up to step 4. There, we first add a a trivial conjunct $\square_a \top$. Then, we are instructed to find $Synth(q \wedge \top)$.

Sub-sub-subgoal 2.1.1: find $Synth(q \wedge \top)$.

We proceed to step 4. There, since there are no \diamond operators in $q \wedge \top$, we return the trivial arrow update model (U_0, o_0) .

Now, in order to find $Synth(\diamond_a q \wedge p)$, we take the trivial arrow update model found in sub-sub-subgoal 2.1.1, and add one extra outcome. Then, we connect this extra outcome to the trivial model by a $\top \rightarrow_a \langle U_0, o_0 \rangle \top$ arrow. The source condition of this arrow is \top because step 4 uses the construction from Lemma 9, and that construction always gives precondition \top . The arrow is for agent a , because we started with a \diamond_a operator. Finally, the target condition is $\langle U_0, o_0 \rangle \top$ because the arrow update that the arrow points to is (U_0, o_0) and the \square_a conjunct was $\square_a \top$. In other words, we obtain the arrow update depicted in Figure 3.a, where the framed state indicates the designated outcome.

Sub-subgoal 2.2: find $Synth(\diamond_a r \wedge p)$.

Replacing the q of $\diamond_a q \wedge p$ for an r does not change the arrow update model that we end up with. So in this sub-subgoal we find the same model as in sub-subgoal 2.1.

Now, in order to find $Synth((\diamond_a q \wedge p) \vee (\diamond_a r \wedge p))$, we need to combine the models found in sub-subgoals 2.1 and 2.2. Since we are working with a disjunction, we combine them as described in Lemma 12. That, is, we take copies of the two (identical) models and add one extra outcome. Then, we add two more arrows: every world that is reachable from the origin world of the model from sub-subgoal 2.1 by $\psi_1 \rightarrow_a \psi_2$, becomes reachable from the extra world by a $\psi_1 \wedge \langle \uparrow \rangle (\diamond_a q \wedge p) \rightarrow_a \psi_2$ arrow. Likewise, every world reachable by $\psi_1 \rightarrow_a \psi_2$ from the origin of the model from sub-subgoal 2.2 becomes reachable from the extra world by $\psi_1 \wedge \neg \langle \uparrow \rangle (\diamond_a q \wedge p) \rightarrow_a \psi_2$. We now get the model depicted in Figure 3.b.

Now, all that is left to do is to combine the arrow update models found in subgoals 1 and 2. The model we obtain is depicted in Figure 3.c, where (U_1, o_1) is the model that we found in subgoal 2. The root of the model is the leftmost outcome. Note that the depth (i.e., the maximum path length) of this arrow update model is 2, just like the depth (i.e., the maximum number of nested \square or \diamond operators) of φ . In general, the depth of the synthesized arrow update model is bounded by that of the formula for which synthesis is performed.




Figure 3: Different stages in the synthesis of $\diamond_a \square_b p \wedge \diamond_b (\diamond_a q \vee \diamond_a r) \wedge \square_b p$

Also, note that the arrow update model that we obtained can quite easily be modified to obtain a smaller model that is still sufficient. In particular:

- the two outcomes that are not reachable from the root can be eliminated,
- the formulas $\langle U_0, o_0 \rangle \top$, $\langle U_1, o_1 \rangle p$, $\top \wedge \langle \uparrow \rangle (\diamond_a q \wedge p)$ and $\top \wedge \neg \langle \uparrow \rangle (\diamond_a q \wedge p)$ can be replaced by the equivalent formulas \top , p , $\diamond_a q \wedge p$ and $\neg(\diamond_a q \wedge p)$, respectively,
- the three leaf outcomes can be merged into one,
- and $\diamond_a q \wedge p \rightarrow_a \top$ and $\neg(\diamond_a q \wedge p) \rightarrow_a \top$ can be merged into one $\top \rightarrow_a \top$ arrow.

With these optimizations, we get the more aesthetically pleasing arrow update depicted in Figure 3.d.

4 Axiomatization

Using the reduction axioms introduced before, we can find an axiomatization for AAUML. Let **AAUML** be the axiomatization shown in Table 2. In this section we show that the axiomatization **AAUML** is sound and complete, and we give some derivable (well-known) axiom schemata.

Prop	all tautologies of propositional logic	
K	$\Box_a(\varphi \rightarrow \psi) \rightarrow (\Box_a\varphi \rightarrow \Box_a\psi)$	
U1	$[U, o]p \leftrightarrow p$	
U2	$[U, o]\neg\varphi \leftrightarrow \neg[U, o]\varphi$	
U3	$[U, o](\varphi \wedge \psi) \leftrightarrow ([U, o]\varphi \wedge [U, o]\psi)$	
U4	$[U, o]\Box_a\varphi \leftrightarrow \bigwedge_{(o, \psi) \rightarrow_a (o', \psi')} (\psi \rightarrow \Box_a(\psi' \rightarrow [U, o']\varphi))$	
A1	$\langle \uparrow \rangle \varphi_0 \leftrightarrow \varphi_0$	where $\varphi_0 \in \mathcal{L}_{pl}$
A2	$\langle \uparrow \rangle (\varphi \vee \psi) \leftrightarrow (\langle \uparrow \rangle \varphi \vee \langle \uparrow \rangle \psi)$	
A3	$\langle \uparrow \rangle (\varphi_0 \wedge \varphi) \leftrightarrow (\varphi_0 \wedge \langle \uparrow \rangle \varphi)$	where $\varphi_0 \in \mathcal{L}_{pl}$
A4	$\langle \uparrow \rangle \bigwedge_{a \in A} (\bigwedge_{\varphi_a \in \Phi_a} \diamond_a \varphi_a \wedge \Box_a \psi_a) \leftrightarrow \bigwedge_{a \in A} \bigwedge_{\varphi_a \in \Phi_a} \diamond_a \langle \uparrow \rangle (\varphi_a \wedge \psi_a)$	
MP	from $\varphi \rightarrow \psi$ and φ infer ψ	
NecK	from φ infer $\Box_a\varphi$	
RE	from $\chi \leftrightarrow \psi$ infer $\varphi[\chi/p] \leftrightarrow \varphi[\psi/p]$	

Table 2: The axiomatization **AAUML** of the logic AAUML

Lemma 17 Axiomatization **AAUML** is sound for the logic AAUML. \dashv

Proof **Prop**, **K**, **MP**, **NecK**, **RE** are known from modal logic, **U1**—**U4** were demonstrated in Section 3.1 and originate in [21], **A1**—**A4** were shown to be valid in Section 3.2. \square

It is important to note that **U1—U4** and **A1—A4** are so-called *reduction axioms* for the operators $[U, o]$ and $\langle \uparrow \rangle$, respectively, as mentioned in the previous section. This means that they are equivalences, where the formula inside the scope of the $[U, o]$ or $\langle \uparrow \rangle$ operator on the left-hand side is more complex than the formulas inside the scope of that operator on the right-hand side.

The derivation rule **RE** is important as our reductions are inside-out, not outside-in. Without it, for example, the validity $[U, o][U, o](p \vee \neg p)$ would not be derivable.

The axioms **A1—A4** could just as well have been formulated with the $[\uparrow]$ dual of the modality $\langle \uparrow \rangle$, e.g., **A2'** $[\uparrow](\varphi \wedge \psi) \leftrightarrow ([\uparrow]\varphi \wedge [\uparrow]\psi)$. We prefer the $\langle \uparrow \rangle$ versions as they match our usage of these axioms in synthesis. Further note that there is no reduction of shape $\langle \uparrow \rangle \neg \varphi \leftrightarrow \dots$. We assume that subformulas bound by $\langle \uparrow \rangle$ are first transformed into disjunctive negation normal form before a further reduction can take place (and again, for this, the derivation rule **RE** is essential).

Lemma 18 Axiomatization **AAUML** is complete for the logic AAUML. \dashv

Proof Let $\varphi \in \mathcal{L}$ be valid. Using an induction argument, we can eliminate all $[U, o]$ and $\langle \uparrow \rangle$ operators from it: φ must be provably equivalent to a formula $\varphi' \in \mathcal{L}_{ml}$. As φ' must also be valid (Theorem 15), φ' is provable in modal logic. From the provable equivalence between φ and φ' and the derivation of φ' we construct a derivation of φ in **AAUML**. \square

We have now shown that:

Theorem 19 Axiomatization **AAUML** is sound and complete for the logic AAUML. \dashv

In the proof system **AAUML**, we do not have necessitation for the $[U, o]$ and $[\uparrow]$ operators. Such necessitation rules are derivable, however.

Proposition 20 The following two rules are derivable in **AAUML**.

- **NecU**: from φ infer $[U, o]\varphi$;
- **NecA**: from φ infer $[\uparrow]\varphi$. \dashv

Proof First, note that the axiom

$$\mathbf{U1'} \quad [U, o]\varphi_0 \leftrightarrow \varphi_0 \quad \text{where } \varphi_0 \in \mathcal{L}_{pl}$$

is derivable, using **Prop**, **U1—U3** and **MP**. It is also convenient to use a variant of **MP** directly on bi-implications, instead of first converting the bi-implication to a single implication.

$$\mathbf{MP'} \quad \text{from } \varphi \leftrightarrow \psi \text{ and } \varphi \text{ infer } \psi, \text{ and from } \varphi \leftrightarrow \psi \text{ and } \psi \text{ infer } \varphi.$$

This **MP'** is, of course, also easily derived. Using **U1'** and **MP'**, we can derive **NecU** in a reasonable number of steps:

1.	φ	premise
2.	$\varphi \rightarrow (\varphi \leftrightarrow \top)$	Prop
3.	$\varphi \leftrightarrow \top$	MP(1,2)
4.	$[U, o]\top \leftrightarrow \top$	U1'
5.	\top	Prop
6.	$[U, o]\top$	MP'(5,4)
7.	$[U, o]\top \leftrightarrow [U, o]\varphi$	RE(3)
8.	$[U, o]\varphi$	MP'(6,7)

A derivation of **NecA** can be found in a similar way. Here, too, it is convenient to first derive an auxiliary axiom.

$$\mathbf{A1}' \quad [\uparrow]\varphi_0 \leftrightarrow \varphi_0 \quad \text{where } \varphi_0 \in \mathcal{L}_{pl}$$

This $[\uparrow]$ -version of **A1** is of course derivable. We can then derive **NecA** analogously to how we derived **NecU**, with the application of **U1'** replaced by **A1'**. \square

The axiomatization **AAUML** is inspired by the somewhat similar axiomatization by Hales of *arbitrary action model logic* [14], although the shape of some axioms and rules is rather different. Axioms in Hales similar to our **U1**–**U4**, for the reduction for action models, are of course taken from action model logic instead. Axioms in Hales similar to our **A1**–**A4**, that are used for the reduction of the quantifier, are taken from refinement modal logic instead. There is nothing ‘of course’ about the latter: Hales’ mixture of AML and RML was, we think, a quite original move. One can also move in the other direction: our axioms **U1**–**U4** allow an alternative axiomatization of refinement modal logic. However, as this seems out of scope, these results are not presented here.

5 Update expressivity

5.1 Expressivity

Recall that we are considering the basic modal logic ML and the update logics PAL, APAL, AML, AAML, AUL, AAUL, AUML, AAUML, and RML, as shown in Figure 1 on page 6. One natural thing to do with such related logics that are all interpreted on a similar class of structures is to compare their power to relate or to distinguish structures from that class. The most straightforward way to make such a comparison is to compare their *expressivity*.

Formally, given languages \mathcal{L}_1 and \mathcal{L}_2 interpreted on certain class of models \mathcal{M} , a language \mathcal{L}_1 is *at least as expressive* as a language \mathcal{L}_2 if for every formula of \mathcal{L}_2 there is an equivalent formula of \mathcal{L}_1 . Having *equal expressivity* or *higher expressivity* (by which we always mean strictly higher expressivity) can be defined from the “at least as expressive” relation in the usual way. If neither language is at least as expressive as the other, we say that they are *incomparable* in expressivity. Informally, if L_1 and L_2 are the logics for

languages \mathcal{L}_1 resp. \mathcal{L}_2 interpreted on \mathcal{M} , we will also use the terminology for expressivity to compare these logics.

In [22] that introduced PAL it was shown that PAL is equally expressive as ML (on the class $\mathcal{S5}$ of relational models, but this does not matter for the reduction), and in [6] that introduced AML it was also shown that AML is equally expressive as ML. It is trivial to show that PAL and AML are at least as expressive as ML, as they extend the logical language. That every formula of PAL or AML is equivalent to a formula in ML, was shown by reduction axioms and rules. Similarly, AUL [20], AUML [21] and AAML [14] were shown to be equally expressive as ML, and therefore also equally expressive as PAL and AML. In [9] it was shown that RML is equally expressive as ML. Here, in Theorem 15 in Section 3, we showed that AAUML is also equally expressive as ML.

The two remaining logics are APAL and AAUL. The logic APAL was shown to be more expressive than ML in [5] and AAUL was shown to be more expressive than ML in [32], wherein it was also shown that APAL and AAUL are incomparable. This means that the expressivity landscape is as shown in Figure 4.




Figure 4: The relative expressivity of the update logics discussed in the paper. Arrows indicate increasing expressivity. Absence of arrows indicates incomparable expressivity.

5.2 Update expressivity hierarchy

There is something a bit strange about this comparison, however. Although AML and PAL have the same expressivity, AML is clearly in some sense more powerful, since action models represent a far larger class of updates than public announcements. In order to capture the sense in which AML is more powerful than PAL, we use the term *update expressivity*. This concept was introduced as *action equivalence* in [35] (and its precursors) and also subsequently used in that sense in [21]. The definitions from [35, 21] do not deal very well with multi-pointed update modalities and with arbitrary update modalities, however, so we use a slightly adapted version.

We recall from the introduction that the updates X we consider are relational model transformers and that such transformations are defined by pairwise relating pointed models:

An update X is a relation between pointed relational models.

In fact, three different things are called update: the *update relation* between pointed models, the *update modality* in a logical language, and, in some sense, the *update object*, often a name, that can be associated with the modality or the relation, such as an arrow update

(U, o) . To simplify the presentation in this section we call the relations X, Y, \dots and the modalities $[X], [Y], \dots$ and we do not consider the update objects separately: we identify them with the update relations.

A relation between pointed relational models can be a one-to-one relation, i.e., a function or a partial function, a one-to-finitely-many relation, and a one-to-infinitely-many relation. For example, it is a function for a pointed arrow update model, a partial function for a public announcement, a one-to-many relation for a multi-pointed arrow update model, and a one-to-infinitely-many relation for an arrow update quantifier. In the first place, one would now like to say that update relations X and Y are the same (are equivalent) if they define the same relation between pointed models, modulo bisimulation. In the second place, we also want to compare an update X that is a partial function, i.e., with a restricted domain, to an update Y that is a total function (or similarly for relations with restricted domains. In that kind of situation one would maybe like to say that updates X and Y are the same if X and Y define the same relation *on the domain of X* : we will then say that X is conditionally equivalent to Y (this relation is asymmetric). Such a requirement seems common practice in dynamic epistemic logic, and it is also respected in [21]. We recall from Section 1 the ‘state eliminating’ public announcement of p (i) and the ‘arrow eliminating’ public announcement of p (ii), originating with [12]: whenever p can be truthfully announced, the pointed relational models resulting from executing (i) and (ii) are bisimilar, as in the example. But when p is false, (ii) can be executed but not (i). So (i) and (ii) are equivalent on condition of the truth of the announcement.

In view of these considerations, we propose the following definition. In the definition, for $X(M, s)$ read $\{(M', s') \mid ((M, s), (M', s')) \in X\}$, and let $dom(X)$ be $\{(M, s) \mid X(M, s) \neq \emptyset\}$, i.e., $dom(X)$ is the domain of X in the standard relational sense.⁵ Recall from Section 2.1 that two sets of pointed models are bisimilar if every pointed model in the first set is bisimilar to one in the second set and vice versa.

Definition 21 (Update equivalence, update expressivity) Given updates X and Y , X is *conditionally update equivalent* to Y , if for all $(M, s) \in dom(X)$, $X(M, s) \Leftrightarrow Y(M, s)$. Further, X is *update equivalent* to Y , if X is conditionally update equivalent to Y , and Y is conditionally update equivalent to X . Update modalities $[X]$ and $[Y]$ are (conditionally) update equivalent, if X and Y are (conditionally) update equivalent.

A language \mathcal{L} is *at least as update expressive as \mathcal{L}'* if for every update modality $[X]$ of \mathcal{L}' there is an update modality $[Y]$ of \mathcal{L} such that X is conditionally update equivalent to Y ; \mathcal{L} is *equally update expressive as \mathcal{L}'* (or ‘as update expressive as’) if \mathcal{L} is at least as update expressive as \mathcal{L}' and \mathcal{L}' is at least as update expressive as \mathcal{L} . \dashv

We define ‘(strictly) more update expressive’ and ‘incomparable in update expressivity’

⁵We recall that given a relational model $M = (S, R, V)$ or an arrow update model $U = (O, RR)$, the sets $S = \mathcal{D}(M)$ resp. $O = \mathcal{D}(U)$ of objects on which the accessibility relations R_a resp. arrow relations RR_a are defined are also called domains. This causes confusion if the arrow update model U is considered as an update relation between sets of pointed relational models, with its domain $dom(U)$ of application defined as above. As the term ‘domain’ is extremely standard for both, we prefer to distinguish them by using different symbols \mathcal{D} and dom , instead of introducing a non-standard term for either one or the other.

as usual. We also extend the usage of ‘update expressive’ to the logics for the languages that we compare. Instead of ‘update equivalent’ we may use ‘equivalent’ if the context is clear. If updates X and Y are update equivalent, then $[X]\varphi \leftrightarrow [Y]\varphi$ is valid. In the other direction, this is not always the case! In Section 7 we give a counterexample.

We should stress that we do not claim that our definition is appropriate for all situations, merely that it gives an accurate view of the strengths of the different logics that we consider in this particular paper.

Let us now fill in the expressivity hierarchy for our target logics. The update expressivity of AUL is higher than that of PAL, and lower than that of AML [20]. The comparison between AML and AUML that we will address in Section 6 is less straightforward than that. In [21] it was shown that AML and AUML have the same update expressivity. That result does not distinguish between single-pointed and multi-pointed action models and arrow update models, however. Here, we therefore provide an alternative proof of their results that makes that distinction. Specifically, we show that the result from [21] only applies to the multi-pointed case, but that single-pointed arrow update models are more update expressive than single-pointed action models.

To formulate such results we need to slightly expand our notation. We recall that in the language \mathcal{L} of AAUML (see Section 2.2) we permit multi-pointed arrow update modalities. If we only allow single-pointed arrow update modalities, let us call the set of validities AAUML₁. For the ‘logic’ in the standard sense of the set of validities this makes no difference, so that out of the context of update expressivity we can continue to also let AAUML represent AAUML₁. Similarly, without the quantifiers, we distinguish AUML₁ from AUML, and we will later also introduce such a notational distinction for action model logics.

Adding quantification increases update expressivity, since the non-quantified logics cannot simulate a one-to-infinity relation. So, for example, APAL is more update expressive than PAL, and AAUML is more update expressive than AUML. (The distinction between single-pointed and multi-pointed is irrelevant here, as quantifying over all single-pointed updates is equivalent to quantifying over all multi-pointed updates.) When comparing the quantified logics among themselves, most pairs are incomparable. These incomparability results are all rather trivial, so we do not prove them here. The only comparable pair is (multi-pointed) AAUML and AAML, which have the same update expressivity because their underlying updates have the same update expressivity (Section 6). In Section 7 we will show that RML is incomparable to the other quantified logics, and in particular that the AAUML and AAML quantifiers are contained in the RML quantifier.

The landscape of update expressivity is therefore as shown in Figure 5.




Figure 5: The relative update expressivity of the update logics discussed in the paper. We assume transitive closure of arrows.

6 Arrow updates versus action models

6.1 Action model logic

Arrow update model logic AUML is equally expressive as action model logic AML and their update expressivity relates in interesting ways. We build upon the results known from [21] but our constructions and proofs are slightly different. First we need to define action models and their execution in relational models. An action model [6] is a structure like a relational model but with a precondition function instead of a valuation function. Executing an action model into a relational model means computing what is known as their restricted modal product. This product encodes the new state of information, after action execution. These are the technicalities:

An *action model* $E = (\mathbf{S}, \mathbf{R}, \text{pre})$ consists of a *domain* \mathbf{S} of *actions* e, f, \dots , an *accessibility function* $\mathbf{R} : \mathbf{A} \rightarrow \mathcal{P}(\mathbf{S} \times \mathbf{S})$, where each \mathbf{R}_a is an accessibility relation, and a *precondition function* $\text{pre} : \mathbf{S} \rightarrow \mathcal{L}$, where \mathcal{L} is a logical language.

Let additional to a pointed action model (E, e) as above a pointed relational model (M, s) be given where $M = (S, R, V)$. Let $M, s \models \text{pre}(s)$. The update $(M \otimes E, (s, e))$ is the pointed relational model where $M \otimes E = (S', R', V')$ such that

$$\begin{aligned}
 S' &= \{(t, f) \mid M, t \models \text{pre}(f)\} \\
 ((t, f), (t', f')) \in R'_a &\text{ iff } (t, t') \in R_a \text{ and } (f, f') \in R_a \\
 (t, f) \in V'(p) &\text{ iff } t \in V(p)
 \end{aligned}$$

Action model modalities $[E, e]$ are interpreted similarly to arrow update modalities but unlike arrow update modalities are partial and not functional. Their execution depends on the truth of the precondition of the actual action (point) e in the actual state s :

$$M, s \models [E, e]\varphi \quad \text{iff} \quad M, s \models \text{pre}(e) \text{ implies } M \otimes E, (s, e) \models \varphi$$

Similarly to arrow update modalities we can conceive a modal logical language with $[E, e]\varphi$ as an inductive language construct, for action models E with finite domains. The logic is called *action model logic* AML. And also similarly we define multi-pointed action models by notational abbreviation, and informally consider such modalities as logical connectives binding formulas. As for AUML (see [21] and Section 4), there is a complete axiomatization, that is a rewrite system allowing to eliminate dynamic modalities [6, 31]. If we further extend the logical language with a quantifier $[\otimes]$ over action models, such that

$$M, s \models [\otimes]\varphi \quad \text{iff} \quad M, s \models [E, e]\varphi \text{ for all action models } (E, e) \text{ satisfying } (*)$$

where $(*)$ requires all preconditions of actions in E to be in \mathcal{L}_{ml} , we get the language and logic of *arbitrary action model logic* AAML. Hales showed in [14] that the $(*)$ requirement can be relaxed, similarly to our Proposition 16. If such a distinction is necessary, the action model logics with only single-pointed action models are AML_1 and AAML_1 .

Example action models that are update equivalent to the example arrow update models of Section 2 are depicted in Figure 6. We also depict their execution. The actions are given their preconditions as names. Note that the pointed relational model resulting from the (second) action of Anne privately learning that p is bisimilar to the four-state model in Section 2 (Figure 2).




Figure 6: Examples of action model execution

6.2 From action models to arrow updates

A given action model can be transformed into an arrow update model by decorating each arrow in the action model with a source condition that is the precondition of the source

action and with a target condition that is the precondition of the target action. That is all. Technically:

Let $E = (\mathbf{S}, \mathbf{R}, \mathbf{pre})$ be given. Arrow update model $U(E) = (O, RR)$ is defined as: $O = \mathbf{S}$, and for all agents a and actions e, e' , $(e, \mathbf{pre}(e)) \rightarrow_a (e', \mathbf{pre}(e'))$ iff $(e, e') \in \mathbf{R}_a$.⁶ We can now show that (E, e) is update equivalent to $(U(E), e)$, on condition of the executability of the action e , i.e., restricted to the denotation of $\mathbf{pre}(e)$.

Proposition 22 ([21]) (E, e) is conditionally update equivalent to $(U(E), e)$. \dashv

Proof Let $M = (S, R, V)$, $M \otimes E = (S', R', V')$ and $M * U(E) = (S'', R'', V'')$. Define the relation \mathfrak{R} by $\mathfrak{R} = \{((s, e), (s, e)) \in S' \times S'' \mid (s, e) \in S'\}$. We will show that if $(M, s) \in \text{dom}(E, e)$, so if $M, s \models \mathbf{pre}(e)$, then \mathfrak{R} is a bisimulation linking $(s, e) \in S'$ to $(s, e) \in S''$.

Take any $((s_1, f_1), (s_1, f_1)) \in \mathfrak{R}$. The **atoms** clause is trivially satisfied as the states s match (and updates do not change facts). We now consider **forth**. Suppose $(s_1, f_1)R'_a(s_2, f_2)$. Given the definition of action model execution, this implies that $(s_1, s_2) \in R_a$, $(f_1, f_2) \in R_a$, $M, s_1 \models \mathbf{pre}(f_1)$ and $M, s_2 \models \mathbf{pre}(f_2)$. Because of how we constructed $U(E)$, it follows from $(f_1, f_2) \in R_a$ that $(f_1, \mathbf{pre}(f_1)) \rightarrow_a (f_2, \mathbf{pre}(f_2))$. Then from $M, s_1 \models \mathbf{pre}(f_1)$ and $M, s_2 \models \mathbf{pre}(f_2)$ it follows that $(s_1, f_1)R''_a(s_2, f_2)$. Also, $((s_2, f_2)(s_2, f_2)) \in \mathfrak{R}$ by the definition of the relation \mathfrak{R} . So **forth** is satisfied. Finally, we consider **back**. Suppose $(s_1, f_1)R''_a(s_2, f_2)$. Then we must have $(f_1, \mathbf{pre}(f_1)) \rightarrow_a (f_2, \mathbf{pre}(f_2))$ with $(f_1, f_2) \in R_a$, since $U(E)$ only contains arrows of that type. It follows that $M, s_1 \models \mathbf{pre}(f_1)$ and $M, s_2 \models \mathbf{pre}(f_2)$, so $(s_1, f_1)R'_a(s_2, f_2)$. Again, obviously, $((s_2, f_2)(s_2, f_2)) \in \mathfrak{R}$. The **back** condition is therefore also satisfied.

We have now shown that \mathfrak{R} is a bisimulation. Since $(s, e)\mathfrak{R}(s, e)$ this implies that $M \otimes E, (s, e) \Leftrightarrow M * U(E), (s, e)$. This holds for every $(M, s) \in \text{dom}(E, e)$, so (E, e) is conditionally update equivalent to $(U(E), e)$. \square

Corollary 23 Let $F \subseteq \mathcal{D}(E)$. Then (E, F) is conditionally update equivalent to $(U(E), F)$. \dashv

From Proposition 22 follows that, for all $\varphi \in \mathcal{L}_{ml}$, $[E, e]\varphi$ is equivalent to $\mathbf{pre}(e) \rightarrow [U(E), e]\varphi$. See also [21][Cor. 3.9]. This can be used as a clause in an inductively defined translation from the language of AUML to the language of AML. In Corollary 23 the condition for update equivalence is $\bigvee_{e \in F} \mathbf{pre}(e)$.

The arrow update models constructed by the above procedure from the action model for Anne reading the letter containing p while Bill may notice her doing so, and for Anne privately learning p , are as follows. Note that they are update equivalent (namely on their domain of execution) to the arrow update models for these actions presented in Section 2. In the figure, by $\varphi \text{ }_{ab} \text{ } \varphi'$ we mean the two arrows $\varphi \rightarrow_a \varphi', \varphi \rightarrow_b \varphi'$. On the left, the dual arrows for b between outcomes have not been labelled. They are as expected: $\top \rightarrow_b \neg p$, $p \rightarrow_b \neg p$, $p \rightarrow_b \top$. (They are update equivalent to the example arrow update models in Section 2.4 and Figure 2.)

⁶In [21], arrows $(e, \top) \rightarrow_a (e', \varphi')$ instead of $(e, \varphi) \rightarrow_a (e', \varphi')$ are stipulated. Both constructions deliver the desired update equivalence.




Figure 7: From action models to arrow updates — example

6.3 From arrow updates to action models

A given arrow update model can be transformed into an update equivalent action model by conditionalizing in each outcome over any possible ‘valuation’ of (any subset of) the source and target conditions of all outcomes. This leads to an exponential blowup. (See [21, Theorem 3.7]. Our construction and subsequent proof are different.) We proceed with the construction.

Let $U = (O, RR)$ be given. Let Φ be the collection of all source and target conditions occurring in U :

$$\Phi = \{\varphi \mid \text{there are } a \in A, \varphi' \in \mathcal{L}, o, o' \in O \text{ s.t. } (o, \varphi) \rightarrow_a (o', \varphi') \text{ or } (o, \varphi') \rightarrow_a (o', \varphi)\}.$$

We consider the formulas in Φ as ‘atomic constituents’ over which we consider ‘valuations’ $v \in 2^\Phi$ (lower case, to distinguish it from the relational model valuation V , upper case). The characteristic formula of a valuation is $\delta_v := \bigwedge_{\varphi \in \Phi} \bar{\varphi}$, where $\bar{\varphi} = \varphi$ if $v(\varphi) = 1$ and $\bar{\varphi} = \neg\varphi$ if $v(\varphi) = 0$. Action model $E(U) = (\mathbf{S}, \mathbf{R}, \mathbf{pre})$ is now such that:

$$\begin{aligned} \mathbf{S} &= O \times 2^\Phi \\ ((o, v), (o', v')) \in \mathbf{R}_a &\text{ iff } \exists \varphi, \varphi' : (o, \varphi) \rightarrow_a (o', \varphi'), v(\varphi) = 1, v'(\varphi') = 1 \\ \mathbf{pre}(o, v) &= \delta_v \end{aligned}$$

Further, given (U, o) , its single point o becomes a set of actions $E(o) := \{o\} \times 2^\Phi$. The corresponding action model $(E(U), E(o))$ is therefore multi-pointed (unless $\Phi = \{\top\}$ or $\Phi = \{\perp\}$). We note that the preconditions of actions need not be consistent formulas, just as source and target conditions of arrows need not be consistent formulas. Our construction is therefore different from that in [21], wherein only $v \in 2^\Phi$ are considered for which δ_v is consistent. That construction is more economical, but computational efficiency is not our goal. We can now show that (U, o) is update equivalent to $(E(U), E(o))$. Note that they both can be executed on the entire domain.

Proposition 24 (U, o) is update equivalent to $(E(U), E(o))$. ⊣

Proof Let $M = (S, R, V)$ and $s \in S$ be given, and let $M * U = (S', R', V')$ and $M \otimes E(U) = (S'', R'', V'')$. We show that for some $(o, v) \in E(o)$, $(M * U, (s, o)) \Leftrightarrow (M \otimes E(U), (s, o, v))$. Define relation \mathfrak{R} as follows:

$$(s, o)\mathfrak{R}(s, o, v) \quad \text{iff} \quad M, s \models \delta_v$$

We show that \mathfrak{R} is a bisimulation.

For **forth**, assume $((s_1, o_1), (s_1, o_1, v_1)) \in \mathfrak{R}$ and $((s_1, o_1), (s_2, o_2)) \in R'_a$. The latter implies that there are $\varphi_1, \varphi_2 \in \mathcal{L}$ such that $(o_1, \varphi_1) \rightarrow_a (o_2, \varphi_2)$, and that $M, s_1 \models \varphi_1$ and $M, s_2 \models \varphi_2$. Choose v_1 and v_2 such that $M, s_1 \models \delta_{v_1}$ and $M, s_2 \models \delta_{v_2}$ (note that v_1 and v_2 exist and are unique). As $M, s_1 \models \varphi_1$ and $M, s_1 \models \delta_{v_1}$ we have $v_1(\varphi_1) = 1$ and as $M, s_2 \models \varphi_2$ and $M, s_2 \models \delta_{v_2}$ we have $v_2(\varphi_2) = 1$. We claim that (s_2, o_2, v_2) is the requested witness to close the **forth** argument. Firstly, $((s_1, o_1, v_1), (s_2, o_2, v_2)) \in R''_a$ because $(s_1, s_2) \in R_a$ and $((o_1, v_1), (o_2, v_2)) \in R_a$, where the latter follows from $(o_1, \varphi_1) \rightarrow_a (o_2, \varphi_2)$, $v_1(\varphi_1) = 1$ and $v_2(\varphi_2) = 1$ (see the definition of $E(U)$ above). Secondly, $((s_2, o_2), (s_2, o_2, v_2)) \in \mathfrak{R}$. Step **back** is very similar; now use that $\text{pre}(o_1, v_1) = \delta_{v_1}$ and that $\text{pre}(o_2, v_2) = \delta_{v_2}$, and observe that $(E(U), E(o))$ can always be executed, it has precondition $\bigvee_{v \in 2^\Phi} \delta_v$, which is equivalent to \top . \square

Corollary 25 Let now $Q \subseteq \mathcal{D}(U)$. Then (U, Q) is update equivalent to $(E(U), E(Q))$. \dashv

As an example we now show the action models constructed by the above procedure from the two example arrow update models of Section 2. The set of source and target condition formulas is $\{p, \neg p, \top\}$. Of their 8 valuations the two non-trivial (and different) valuations are characterized by p and $\neg p$. These formulas are also the action preconditions in the action points of the resulting action models. The reader may observe that these action models are, again, update equivalent to their ‘original’ action models at the start of this section.




Figure 8: From arrow updates to action models — example

6.4 Relative update expressivity

We are now prepared to harvest the update expressivity results. First, let us show that AUML_1 is more update expressive than AML_1 .

Proposition 26 AUML_1 is more update expressive than AML_1 . \dashv

Proof From Proposition 22 follows that AUML_1 is at least as update expressive as AML_1 . To show that the inclusion is strict, we need to show that for some (U, o) , there is no (E, e) that induces the same relation.

Let U be the arrow update model with a single outcome o , and a single arrow $(o, p) \rightarrow_a (o, \top)$. Suppose towards a contradiction that there is a single-pointed action model (E, e)

such that $(M * U, (s, o)) \Leftrightarrow (M \otimes E, (s, e))$ for every (M, s) . Then, in particular, (E, e) must be executable everywhere, so $\mathbf{pre}(e)$ is equivalent to \top . Furthermore, if $M, s \models p \wedge \diamond_a \top$, then $(M * U, (s, o))$ has at least one a -successor (s', o) (note that $\diamond_a \top$ enforces that s has an a -successor s' , and the arrow $s \rightarrow_a s'$ satisfies p at the source and \top at the target). By assumption, $(M \otimes E, (s, e))$ is bisimilar to $(M * U, (s, o))$, so it has an a -successor (s', e') . This implies that $eR_a e'$. Now, consider a different model (M', s) such that (i) $M', s \models \neg p$ and (ii) s has an a -successor s' such that $M', s' \models \mathbf{pre}(e')$. Then $(M' \otimes E, (s, e))$ has an a -successor, but $(M' * U, (s, o))$ has no a -successor, because p is false in state s . This contradicts the assumption that $(M * U, (s, o)) \Leftrightarrow (M \otimes E, (s, e))$ for every (M, s) . \square

Proposition 27 AUML is as update expressive as AML. \dashv

Proof From Corollary 23 follows that AUML is at least as update expressive as AML. From Corollary 25 follows that AML is at least as update expressive as AUML. \square

It is obvious that AML is more update expressive than AML_1 , and that AUML is more update expressive than AUML_1 .

6.5 Applications illustrating the succinctness of arrow updates

In this section we give some application areas for the modelling of information change, where arrow updates are more succinct than corresponding (i.e., update equivalent) action models.

Lying You are lying if you say that something is true while you believe that it is false — and with the intention for the addressee(s) to believe that it is true. In the setting of public announcement logic a lie is a public announcement that is false. This is then contrasted to the (usual) public announcement that is true. Both are combined in the announcement that has no relation to its truth. This is known as the conscious update ([12], see the introduction) or, in a setting where lying is also distinguished, as the manipulative update [34]. The arrow update for the conscious/manipulative update of φ is the singleton arrow update model with arrows

$$(o, \top) \rightarrow_a (o, \varphi)$$

for all agents. Lying as such, wherein φ is required to be false, is not an arrow update as arrow update models have no preconditions, but such executability preconditions can be simulated as antecedents of logical implications.

A problem with the manipulative update is that an agent who already believes the opposite of the lie, believes everything after incorporating the lie into her beliefs (believing a contradiction comes at that price). This is because the accessibility of that agent becomes empty as a result of the update. A solution to that is the *cautious update* that is also known as lying to *sceptical* agents [23, 20, 27]: the agent only updates her beliefs if the new information is consistent with her current beliefs. The arrow update for the *sceptical*

update of φ (again, we cannot model sceptical *lying* as this requires φ to be false) is a singleton arrow update model with arrows

$$\begin{aligned} (o, \Diamond_a \varphi) &\rightarrow_a (o, \varphi) \\ (o, \Box_a \neg \varphi) &\rightarrow_a (o, \top) \end{aligned}$$

for all agents [20].

Given a group of agents, some may believe the announcement, and others not. The arrow update modelling allows for this. However, if we make an action model for announcements to sceptical agents, we need to distinguish all combinations explicitly (we used a similar construction to get action model $E(U)$ from arrow update U , above, in order to prove Prop. 24). For example, for two agents a and b , the action model consists of eight actions, with preconditions and accessibility relations as follows. In Figure 9 we ‘name’ the actions with their preconditions. To simplify the visualization, we do not label arrows with a and b : solid arrows are for a and dashed arrows for b . We also assume transitive closure of accessibility.




Figure 9: Action model for lying

For n agents there are $O(2^n)$ actions in the action model. In [20] this example is treated in greater detail, and also other, similar, examples are shown for which arrow updates are shown to be more succinct (exponentially smaller).

Attentive announcements Another example where action models are exponentially bigger than arrow updates is that of the *attention-based announcements* of [8]. This work presents a logic of announcements that are only ‘heard’ (received) by agents paying attention to it, paying attention to the announcer, so to speak. Such announcements are modelled employing an auxiliary set of designated ‘attention (propositional) variables’ h_a

expressing that agent a pays attention. The corresponding arrow update model has domain $\{o, o'\}$, both outcomes designated, and with arrows

$$\begin{aligned} (o, h_a) &\rightarrow_a (o, \varphi) \\ (o, \neg h_a) &\rightarrow_a (o', \top) \\ (o', \top) &\rightarrow_a (o', \top) \end{aligned}$$

for all agents. It cannot be modelled with a (singleton) [20] arrow update, the resulting relational model is typically larger than then model before the update, as the agents not paying attention believe that no announcement was made, and thus reason about the structure of the entire initial model. Incorporating the announcements depends on h_a being true or false, just as for sceptical announcements it depends on $\Diamond_a \varphi$ being true or false. So, this is similar. But not entirely so, because an agent not paying attention is, so to speak, unconscious of the announcement, and thus believes that the (entire) original model still encodes her beliefs), whereas an agent believing the opposite of the announcement ‘knows’ that if she where to have found the announcement believable, she would have changed her beliefs. So these are different parts of the same model, it is a mere restriction of the accessibility relation. Again, for attentive announcements, a corresponding action model is of exponential size, as any subset of agents may or may not be paying attention. See [8].

Comparative size of action models and arrow updates In general, if the observational powers of *all* agents are commonly known to be partial, then we can expect arrow updates for such dynamic phenomena to be exponentially smaller than corresponding action models. This was the case for announcements to sceptical lying and for attention-based announcements, and also for: agents making broadcasts (to all agents), agents seeing each other depending on their orientation, partial networks representing agents with neighbours or friends, etc. On the other hand, dynamic phenomena where all agents observe (some, few) designated agents have similarly-sized arrow updates and action models, such as: the private announcement to an individual agent or a subgroup of agents, and gossip scenarios where two agents call each other in order to exchange secrets, and where this call may be partially observed by all other agents. We do not know of scenarios where action models are more succinct than arrow updates.

7 Arbitrary arrow updates versus refinements

7.1 Refinement modal logic

We now compare the arbitrary arrow update modality of AAUML to the *refinement* quantifier of refinement modal logic RML [9]. Let us first be precise about its syntax and semantics.

We recall the definition of bisimulation in Section 2.1. If **atoms** and **back** hold, we call the relation a *refinement* (and dually, if **atoms** and **forth** hold, we call the relation

a *simulation*). In [9] such a refinement relation is considered for any subset of the set of agents and defined as follows:

A relation \mathfrak{R}_B that satisfies **atoms**, **back- a** , and **forth- a** for every $a \in A \setminus B$, and that satisfies **atoms**, and **back- b** for every $b \in B$, is a *B-refinement*, we say that (M', s') *refines* (M, s) for group of agents B , and we write $(M, s) \succeq_B (M', s')$. An *A-refinement* we call a *refinement* (clearly any B -refinement is also an A -refinement and thus a ‘refinement’ plain and simple), and $(M, s) \succeq_A (M', s')$ is denoted $(M, s) \succeq (M', s')$. With this relation comes a corresponding modality in the obvious way. Let $B \subseteq A$, and (M, s) and φ given, then

$$M, s \models [\succeq]_B \varphi \quad \text{iff} \quad M', s' \models \varphi \text{ for all } (M', s') \text{ such that } (M, s) \succeq_B (M', s').$$

We continue with the comparison of the refinement quantifier to the other quantifiers. We first focus on the refinement relation \succeq for the set of all agents, to which corresponds the $[\succeq]$ modality. Consider three different ways to define quantification in information changing modal logics. We formulate them suggestively so that their correspondences stand out, where we recall Proposition 16 that the restrictions on source and target conditions need not be met when interpreting $[\uparrow]$, and similarly, [14] showed that the restrictions on action preconditions need not be met when interpreting $[\otimes]$.

$$\begin{aligned} M, s \models [\uparrow] \varphi & \quad \text{iff} \quad M, s \models [U, o] \varphi \text{ for all arrow update models } (U, o) \\ M, s \models [\succeq] \varphi & \quad \text{iff} \quad M', s' \models \varphi \text{ for all refinements } (M', s') \\ M, s \models [\otimes] \varphi & \quad \text{iff} \quad M', s' \models [E, e] \varphi \text{ for all action models } (E, e) \end{aligned}$$

Theorem 28 Let $\varphi \in \mathcal{L}_{ml}$. Then $[\uparrow] \varphi$, $[\succeq] \varphi$, $[\otimes] \varphi$ are pairwise equivalent. ⊣

Proof

- The equivalence of $[\succeq] \varphi$ to $[\otimes] \varphi$ was shown in [14].
- To show that $[\otimes] \varphi$ is equivalent to $[\uparrow] \varphi$, we use the semantics of these modalities. Let us do this for the diamond version. Both directions of the equivalence need to be shown.

First, suppose that

$$M, s \models \langle \otimes \rangle \varphi.$$

According to the semantics of $\langle \otimes \rangle$ this is equivalent to

$$\exists(E, e) : M, s \models \langle E, e \rangle \varphi.$$

By Proposition 22 that, in turn, is equivalent to

$$\exists(U(E), e) : M, s \models \mathbf{pre}(e) \text{ and } M, s \models \langle U(E), e \rangle \varphi.$$

This implies that, in particular,

$$\exists(U(E), e) : M, s \models \langle U(E), e \rangle \varphi$$

and therefore, by the semantics of $\langle \uparrow \rangle$, that

$$M, s \models \langle \uparrow \rangle \varphi.$$

For the other direction, suppose that

$$M, s \models \langle \uparrow \rangle \varphi.$$

According to the semantics of $\langle \uparrow \rangle$ this is equivalent to

$$\exists(U, o) : M, s \models \langle U, o \rangle \varphi.$$

By Proposition 24 that, in turn, is equivalent to

$$\exists(E(U), E(o)) : M, s \models \langle E(U), E(o) \rangle \varphi.$$

In particular, this implies that for some $e \in E(o)$ we have

$$\exists(E(U), e) : M, s \models \langle E(U), e \rangle \varphi$$

and therefore, by the semantics of $\langle \otimes \rangle$, that

$$M, s \models \langle \otimes \rangle \varphi.$$

- To show that $[\uparrow]\varphi$ is equivalent to $[\succeq]\varphi$ we use the previous two equivalences. □

The theorem is formulated to make the correspondence between the three quantifiers stand out. Alternatively, we can have an inductively defined translation between the language \mathcal{L} (of AAUML) and the language of arbitrary action model logic AAML that is compositional to the extent that arrow update quantifiers are translated into action model quantifiers (Theorem 28) and arrow update models into action models (Proposition 24), and vice versa (Proposition 22).

7.2 Update expressivity

Considering that $[\otimes]\varphi$, $[\uparrow]\varphi$ and $[\succeq]\varphi$ are equivalent for basic modal formulas φ (Theorem 28), and that $[\otimes]$ and $[\uparrow]$ have the same update expressivity (Section 5), one might expect all three logics AAML, AAUML, and RML to have the same update expressivity. This, however, is not so, because $[\otimes]$ and $[\uparrow]$ are finitary quantifiers — they quantify over, respectively, *finite* action models and over *finite* arrow update models — whereas refinements can be infinitary.

For one example, consider the relational model N consisting of all valuations, with the universal relation on that domain for all agents, and any state t in that domain. Clearly, the restriction of N to the singleton model consisting of t (wherein the agents have common knowledge of the valuation in t) is a refinement of (N, t) . It can be obtained by successively announcing the value of each of the infinite number of atoms. However, it cannot be obtained by a single announcement (or, equivalently, by any finite sequence of those).

For another example, consider the models M , with as single state s_0 , and M' , with s_0 as its leftmost state, in Figure 10. The pointed model (M', s_0) is a refinement of (M, s_0) . But M' contains infinitely many states that are not bisimilar to one another. Furthermore,

every arrow update model U in the logical language is finite, so every product of U with M is finite (and therefore contains finitely many non-bisimilar states). As a result, there is no (U, o) such that (M', s_0) is bisimilar to $(M * U, (s, o))$.




Figure 10: Arbitrary arrow updates and refinements are incomparable

It follows that arbitrary arrow updates are *not* at least as update expressive as refinements. But it also follows that refinements are not at least as update expressive as arbitrary arrow updates, since you cannot choose to exclude the above model (M', s_0) when performing a refinement in (M, s_0) .

Proposition 29 RML and AAUML are incomparable in update expressivity. \dashv

And therefore RML and AAML are also incomparable. The reason that $[\uparrow]\varphi$ and $[\succeq]\varphi$ (and $[\otimes]\varphi$) are nonetheless equivalent is that while there is no (U, o) such that $(M * U, (s_0, o))$ is bisimilar to (M', s_0) , it is the case that for every n there is an (U_n, o_n) such that $(M * U_n, (s_0, o_n))$ is n -bisimilar to (M', s_0) . Since every formula in the languages under consideration is of finite depth, such finite approximations of M' suffice.

Finally, we should note that the incomparability already applies to the language for RML with only the $[\succeq]$ modality. The language above, as in [9], has $[\succeq]_B$ modalities for any subgroup $B \subseteq A$, meaning that, modulo bisimulation, only arrows in B are removed from a relational model. Similarly to the argument above it follows that this would only further increase expressivity.

8 Conclusions and further research

Conclusions We presented *arbitrary arrow update model logic* (AAUML). We provided an axiomatization of AAUML, which also demonstrates that AAUML is decidable and equally expressive as multi-agent modal logic. We established arrow update model synthesis for AAUML. We determined the update expressivity hierarchy including AAUML and many other update logics, including other arrow update logics, action model logics, and refinement modal logic.

Further research on B -restricted arrow update synthesis Let B be any subset of the set of all agents. Building upon the B -refinements of [9] and motivated by a similar

approach used in [14], a variant of the synthesis problem for AAUML is to consider *B-restricted* arrow update models. Roughly speaking, a *B-restricted* arrow update model represents an event where only the agents in *B* can gain more factual information, while the agents outside *B* remain at least as uncertain as they were before the event. The *B-restricted* synthesis problem can be solved in a very similar way to the unrestricted problem that we presented in this paper.

Similarly to how arrow update models have the same update expressivity as action models and refinements, *B-restricted* arrow update models have the same update expressivity as *B-restricted* action models, and *B-refinements* have larger update expressivity.

Formally introducing *B-restricted* arrow update models, and showing that the results apply there as well, would require a lot of complicate notation and several complex definitions. So for the sake of readability we did not include them in this paper.

Knowledge and belief Arrow updates result in changes of knowledge and belief. Of particular interest are therefore updates that preserve the $\mathcal{S}5$ or $\mathcal{KD}45$ nature of relational models. It is unclear how to enforce such preservation semantically, as discussed in Section 2.4. Such issues need to be resolved in order to find an axiomatization for arrow update logic, or arrow update model logic, or quantified versions of these, for the class $\mathcal{S}5$ or $\mathcal{KD}45$. These well-known problems [20, 32] are related to similar issues for refinement modal logic, as refinement also is relational change. Non-trivial $\mathcal{S}5$ and $\mathcal{KD}45$ versions of refinement modal logic have been proposed in work by Hales *et al.* [16, 17], and particular mention deserves Hales’ Ph.D. thesis [15, Section 9.3, Section 9.4] wherein the axiomatization AAUML is adapted to the classes $\mathcal{KD}45$ and $\mathcal{S}5$, respectively. These diverse results might inspire similar solutions for truly ‘epistemic’ arrow update logics.

Work in progress on the complexity of synthesis We have shown that it is possible to perform synthesis for AAUML, and described an algorithm that does this synthesis. We have not, however, discussed the *computational complexity* of that algorithm. The complexity is non-elementary. The non-elementary blowup occurs when $\langle \uparrow \rangle$ operators are nested. In general, if φ is a basic modal formula, then the procedure outlined in this paper allows us to find a basic modal formula φ' that is equivalent to $\langle \uparrow \rangle \varphi$, but the size of this φ' is, in the worst case, exponential in the size of φ . As a result, a formula ψ containing n nested $\langle \uparrow \rangle$ operators can be translated to an equivalent formula ψ' of modal logic using this method, but both the formula ψ' and the arrow update model synthesized for the outermost $\langle \uparrow \rangle$ operator will have experienced exponential blowup n times, resulting in a non-elementary blowup overall. We suspect that this is unavoidable, i.e., that the difficulty of the synthesis problem is non-elementary. We do not, for now, have a hardness proof, however.

References

- [1] C. Areces, R. Fervari, and G. Hoffmann. Moving arrows and four model checking results. In *Proc. of 19th WoLLIC*, pages 142–153. Springer, 2012. LNCS 7456.
- [2] G. Aucher. Characterizing updates in dynamic epistemic logic. In *Proceedings of Twelfth KR*. AAAI Press, 2010.
- [3] G. Aucher. DEL-sequents for progression. *Journal of Applied Non-Classical Logics*, 21(3-4):289–321, 2011.
- [4] G. Aucher, P. Balbiani, L. Fariñas del Cerro, and A. Herzig. Global and local graph modifiers. *Electr. Notes Theor. Comput. Sci.*, 231:293–307, 2009.
- [5] P. Balbiani, A. Baltag, H. van Ditmarsch, A. Herzig, T. Hoshi, and T. De Lima. ‘Knowable’ as ‘known after an announcement’. *Review of Symbolic Logic*, 1(3):305–334, 2008.
- [6] A. Baltag, L.S. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. In *Proc. of 7th TARK*, pages 43–56. Morgan Kaufmann, 1998.
- [7] P. Blackburn, J. van Benthem, and F. Wolter, editors. *Handbook of Modal Logic*. Elsevier, 2006.
- [8] T. Bolander, H. van Ditmarsch, A. Herzig, E. Lorini, P. Pardo, and F. Schwarzentruber. Announcements to attentive agents. *Journal of Logic, Language and Information*, 25(1):1–35, 2016.
- [9] L. Bozzelli, H. van Ditmarsch, T. French, J. Hales, and S. Pinchinat. Refinement modal logic. *Information and Computation*, 239:303–339, 2014.
- [10] T. French, J. Hales, and E. Tay. A composable language for action models. In R. Goré, B.P. Kooi, and A. Kurucz, editors, *Advances in Modal Logic 10*, pages 197–216. College Publications, 2014.
- [11] T. French and H. van Ditmarsch. Undecidability for arbitrary public announcement logic. In *Advances in Modal Logic 7*, pages 23–42, London, 2008. College Publications.
- [12] J.D. Gerbrandy. *Bisimulations on Planet Kripke*. PhD thesis, University of Amsterdam, 1999. ILLC Dissertation Series DS-1999-01.
- [13] J.D. Gerbrandy and W. Groeneveld. Reasoning about information change. *Journal of Logic, Language, and Information*, 6:147–169, 1997.
- [14] J. Hales. Arbitrary action model logic and action model synthesis. In *Proc. of 28th LICS*, pages 253–262. IEEE, 2013.

- [15] J. Hales. *Quantifying over epistemic updates*. PhD thesis, School of Computer Science & Software Engineering, University of Western Australia, 2016. <https://research-repository.uwa.edu.au/en/publications/quantifying-over-epistemic-updates>.
- [16] J. Hales, T. French, and R. Davies. Refinement quantified logics of knowledge. *Electr. Notes Theor. Comput. Sci.*, 278:85–98, 2011.
- [17] J. Hales, T. French, and R. Davies. Refinement quantified logics of knowledge and belief for multiple agents. In *Advances in Modal Logic 9*, pages 317–338. College Publications, 2012.
- [18] B. Kooi. *Knowledge, Chance, and Change*. PhD thesis, University of Groningen, 2003. ILLC Dissertation Series DS-2003-01.
- [19] B. Kooi. Expressivity and completeness for public update logics via reduction axioms. *Journal of Applied Non-Classical Logics*, 17(2):231–254, 2007.
- [20] B. Kooi and B. Renne. Arrow update logic. *Review of Symbolic Logic*, 4(4):536–559, 2011.
- [21] B. Kooi and B. Renne. Generalized arrow update logic. In *Proc. of 13th TARK*, pages 205–211, 2011. Poster presentation.
- [22] J.A. Plaza. Logics of public communications. In *Proc. of the 4th ISMIS*, pages 201–216. Oak Ridge National Laboratory, 1989.
- [23] D. Steiner. A system for consistency preserving belief change. In *Proc. of the ESSLLI Workshop on Rationality and Knowledge*, pages 133–144, 2006.
- [24] J. van Benthem. An essay on sabotage and obstruction. In *Mechanizing Mathematical Reasoning*, volume 2605 of *LNCS 2605*, pages 268–276. Springer, 2005.
- [25] J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. *Information and Computation*, 204(11):1620–1662, 2006.
- [26] H. van Ditmarsch. *Knowledge games*. PhD thesis, University of Groningen, 2000. ILLC Dissertation Series DS-2000-06.
- [27] H. van Ditmarsch. Dynamics of lying. *Synthese*, 191(5):745–777, 2014.
- [28] H. van Ditmarsch, J.Y. Halpern, W. van der Hoek, and B. Kooi, editors. *Handbook of epistemic logic*. College Publications, 2015.
- [29] H. van Ditmarsch and B. Kooi. Semantic results for ontic and epistemic change. In *Proc. of 7th LOFT*, Texts in Logic and Games 3, pages 87–117. Amsterdam University Press, 2008.

- [30] H. van Ditmarsch, W. van der Hoek, and B. Kooi. Dynamic epistemic logic with assignment. In *Proc. of 4th AAMAS*, pages 141–148. ACM, 2005.
- [31] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer, 2008.
- [32] H. van Ditmarsch, W. van der Hoek, B. Kooi, and L.B. Kuijer. Arbitrary arrow update logic. *Artif. Intell.*, 242:80–106, 2017.
- [33] H. van Ditmarsch, W. van der Hoek, and L.B. Kuijer. The undecidability of arbitrary arrow update logic. *Theor. Comput. Sci.*, 693:1–12, 2017.
- [34] H. van Ditmarsch, J. van Eijck, F. Sietsma, and Y. Wang. On the logic of lying. In *Games, Actions and Social Software*, LNCS 7010, pages 41–72. Springer, 2012.
- [35] J. van Eijck, J. Ruan, and T. Sadzik. Action emulation. *Synthese*, 185(1):131–151, 2012.
- [36] J. van Eijck, F. Sietsma, and Y. Wang. Composing models. *Journal of Applied Non-Classical Logics*, 21(3-4):397–425, 2011.