

Epistemic Planning with Attention as a Bounded Resource

Gaia Belardinelli and Rasmus K. Rendsvig

Center for Information and Bubble Studies, University of Copenhagen
{belardinelli,rasmus}@hum.ku.dk

Abstract. Where information grows abundant, attention becomes a scarce resource. As a result, agents must plan wisely how to allocate their attention in order to achieve epistemic efficiency. Here, we present a framework for multi-agent epistemic planning with attention, based on Dynamic Epistemic Logic (DEL, a powerful formalism for epistemic planning). We identify the framework as a fragment of standard DEL, and consider its plan existence problem. While in the general case undecidable, we show that when attention is required for learning, all instances of the problem are decidable.

1 Introduction

The development of autonomous agents is central in artificial intelligence. A core feature of autonomous agents is their ability to exhibit goal-directed behaviour, i.e. to commit to goals and generate plans to achieve them. *Epistemic planning* [6] focuses on domains where agents’ plans must take into account their own capabilities and knowledge, as well as knowledge about other agents’ knowledge. For example, in an *epistemic planning problem*, agent a may have as goal “ a knows the truth-value of φ , while b does not know that a knows it”. To achieve this goal, agent a may need to reason about what it can do to learn φ , and about what b may learn about φ from a ’s actions. *Dynamic Epistemic Logic* (DEL) offers a highly expressive basis for epistemic planning, allowing e.g. nondeterminism, partial observability and arbitrary levels of higher-order reasoning [6].

Attention is relevant to autonomous agents that labour in information-rich environments. As H. Simon wrote in [18]: “In an information-rich world, the wealth of information means a dearth of something else: a scarcity of whatever it is that information consumes. What information consumes is rather obvious: it consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention, and a need to allocate that attention efficiently (...).” Attention is thus a bounded resource, crucial to agents that must process information to achieve their goals. In such cases, plans must factor in attention and be attention-feasible: If agents undertake actions that require more attention than is available, they will fail. Autonomous agents in multi-agent systems must then turn to epistemic planning with attention as a bounded resource.

However, DEL does not assume resource bounded agency, and so DEL-based epistemic planning does not bound the attention agents may consume to achieve

epistemic goals. Therefore, we propose an alteration of DEL that incorporates attention as a bounded resource in epistemic planning.

In the proposed model, attention is used in learning. The model is of the ‘DEL-type’ by its ‘state-action-product’ format, but with a twist. *Attention states* portray static snapshots of agents’ information, with agents’ remaining attention encoded propositionally. Attention states are updated with *attention actions*, that have two elements each: an *action model* interpreted as an *information source*, and, for each agent, a *question* asked of the information source (this is the twist). The question, a formula, is what is being paid attention to. Jointly, action model and questions invoke attention cost and information change, calculated by taking a state-action product. Sec. 2 introduces and exemplifies this.

In Sec. 3, we relate to DEL in three ways. We show that attention states may be recast as DEL epistemic states and *vice versa*; that every DEL action *without* postconditions may be emulated by an attention action, but not *vice versa*; and that every attention action may be emulated by a DEL action *with* postconditions, but not *vice versa*. We use these results in Sec. 4.

In Sec. 4, we turn to planning and our main results: The *plan existence problem* for epistemic planning with attention is in general undecidable, but is decidable for *No Free Lunch* (NFL) actions. In NFL actions, complex learning requires questioning, and all non-trivial questions have non-zero attention cost. NFL actions are of special interest as they enforce that attention expenditure is required for active learning. The first result is a corollary to the undecidability of DEL-based planning without postconditions [7] and results in Sec. 3. The second is a consequence of bounded attention. When attention is exhausted, the NFL actions supply a restricted form of background information, ensuring that eventually no further change occurs.

Sec. 5 concludes, discussing how our decidability result differs from those established in DEL-based epistemic planning, and future research questions.

To the best of the authors’ knowledge, the paper’s focus on attention as a bounded resource in epistemic planning is novel. Anyway, some works in the literature relate to different aspects of the present proposal. [3] proposes a model where explicit beliefs (interpreted as the focus of agents attention) may be boundedly extended by perception or inference from memory, using a DEL variant. [19] use DEL to model cognitively costly inference for single agents, using impossible worlds. Both proposals do not relate to planning. Works on time-bounded reasoning, where inference requires time, may be interpreted in terms of attention [1]. These models do not use DEL. Some DEL papers [9, 11] model agents that pay attention or not, affecting whether they learn, and [15] discusses *joint attentional states*. Neither work captures attention as a *resource*. Some works on awareness draw parallels with attention [17], but without focusing on resource boundedness. More peripherally lies the field of *attention economics* [14] which studies attention allocation in markets, but it does not represent agents’ epistemic states, contra DEL.

In the following, the proofs of some propositions will only be sketched, due to page limits. Their full version can be found in arxiv.org/abs/2105.09976.

2 Attention in States and Actions

Attention may be conceptualized in several ways. It may be equated with the time, memory or ‘mental energy’ required to learn some proposition. To accommodate this broad scope, we represent agents’ attention simply as a numerical value with attention costs relative to agent, question and context (event).

2.1 Language

Throughout, let $I \neq \emptyset$ be a finite set of *agents*, let $N \in \mathbb{N}$ be an *attention bound*. We use a classic epistemic logical language, but with a bi-partitioned set of atoms:

Let Φ be a countable set of *proposition atoms* and let $\Psi = \{(\alpha_i < n), (\alpha_i = n) : n \leq N, i \in I\}$ be a set of *attention atoms*. Let the full set of *atoms* be their disjoint union $At = \Phi \uplus \Psi$. With $i \in I, p \in At$, define the language \mathcal{L} by

$$\varphi ::= \top \mid p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_i\varphi$$

The formula $K_i\varphi$ states that agent i knows that φ and the attention atoms $(\alpha_i < n)$ and $(\alpha_i = n)$ state that i has, respectively, strictly less than n or exactly n , attention left, where n is a natural number.

As abbreviations, let $(\alpha_i > n) := \neg(\alpha_i < n) \wedge \neg(\alpha_i = n)$ and $(\alpha_i \geq n) := \neg(\alpha_i < n)$. In figures, we denote $\neg\varphi$ by $\bar{\varphi}$.

2.2 Attention States

We use special cases of Kripke models, augmented with maps to quantify the attention span of each agent at each world (examples are in Sec. 2.3). Sec. 3 shows them recast as Kripke models with valuations suitable for attention.

Definition 1. An attention model is an $M = (W, R, V, A)$ where $W \neq \emptyset$ is a finite set of worlds; $R : I \rightarrow \mathcal{P}(W^2)$ assigns each $i \in I$ an equivalence relation R_i ; $V : At \rightarrow \mathcal{P}(W)$ is a valuation; and $A : I \times W \rightarrow \{0, \dots, N\}$ is an attention resource function satisfying that wR_iv implies $A_i(w) = A_i(v)$, for all $i \in I$. For $w \in W$, (M, w) is an *attention state*, with w the *actual world*.

Relations R_i are taken to capture the indistinguishability of worlds for agent i , and are therefore assumed to be equivalence relations, as standard in epistemic logic [16]. The restriction on A_i ensures that i knows their own attention span.

Definition 2. Let $(M, w) = ((W, R, V, A), w)$ be an attention state. For $i \in I, n \in \mathbb{N}, p \in \Phi, (\alpha_i < n), (\alpha_i = n) \in \Psi$, satisfaction of \mathcal{L} formulas is given by

$$\begin{aligned} M, w \models p & \quad \text{iff } w \in V(p) \text{ for all } p \in \Phi, \\ M, w \models (\alpha_i = n) & \quad \text{iff } A_i(w) = n, \\ M, w \models (\alpha_i < n) & \quad \text{iff } A_i(w) < n, \\ M, w \models K_i\varphi & \quad \text{iff } wR_iv \text{ implies } M, v \models \varphi \end{aligned}$$

with standard clauses for $\top, \neg\varphi$ and $\varphi \wedge \psi$, where $\varphi, \psi \in \mathcal{L}$. If for all $(M, w), (M, w) \models \varphi \rightarrow \psi$, we write $\varphi \models \psi$.

Our results rely on establishing equivalence of models, for which we need to introduce *attention bisimulation*. Attention bisimulation is defined akin to bisimulation for Kripke models, but factoring attention into the Atoms clause.

Definition 3. Any two attention states $M = ((W, R, V, A), w)$ and $M' = ((W', R', V', A'), w')$ are (attention) bisimilar (written $(M, w) \simeq (M', w')$) if there exists a relation $Z \subseteq W \times W'$ such that wZw' and for all $v \in W$, $v' \in W'$, if vZv' , then for all $p \in \Phi, i \in I$,
Atoms $v \in V(p) \Leftrightarrow v' \in V'(p)$ and $A_i(v) = A'_i(v')$;
Forth if vR_iu , then for some $u' \in W'$, $v'R'_iu'$ and uZu' ;
Back if $v'R'_iu'$, then for some $u \in W$, vR_iu and $u'Zu$.

Bisimulation between attention states implies modal equivalence:

Proposition 4. If attention states (M, w) and (M', w') are bisimilar, then for every $\varphi \in \mathcal{L}$, $(M, w) \models \varphi$ iff $(M', w') \models \varphi$.

The proof follows Prop. 13 (Sec. 3), of which it makes use.

Finally, in showing our main decidability result, Theorem 25, we use bisimulation contractions and the following Lemma.

Definition 5. Let $(M, w) = ((W, R, V, A), w)$ be an attention model. The bisimulation contraction of (M, w) is the attention model $(M, w)_{\simeq} = ((W', R', V', A'), [w])$ with $W' = \{[w] : w \in W\}$ for $[w] = \{v \in M : (M, w) \simeq (M, v)\}$; for all $i \in I$, $R'_i = \{([w], [v]) : \exists w' \in [w], \exists v' \in [v] \text{ with } w'R_iv'\}$; $V' : \Phi \rightarrow \mathcal{P}(W)$ with $V'(p) = \{[w] \in W : w \in V(p)\}$ for all $p \in \Phi$; and $A'_i([w]) = A_i(w)$, for all $w \in W$.

Lemma 6. For any attention state (M, w) , for all $\varphi \in \mathcal{L}$, $(M, w)_{\simeq} \models \varphi$ iff $(M, w) \models \varphi$.

Proof (Proof sketch). $(M, w)_{\simeq}$ is bi-similar to (M, w) , witnessed by $Z = \{(w, [w]) : w \in W\}$. The conclusion then follows by Prop. 4.

2.3 Attention Actions

To perform actions on attention states, we use an augmented version of action models (without postconditions), well-known from DEL [4], plus a map capturing what question (formula) each agent pays attention to.

Definition 7. An attention action model is a tuple $\mathcal{A} = (E, Q, Q^*, pre, c)$ where $E \neq \emptyset$ is a finite set of events; $Q : I \rightarrow \mathcal{P}(E^2)$ and $Q^* : I \rightarrow \mathcal{P}(E^2)$ assign each $i \in I$ equivalence relations Q_i and Q_i^* ; $pre : E \rightarrow \mathcal{L}$ assigns a precondition to each event; $c : I \times \mathcal{L} \times E \rightarrow \mathbb{N}$ is a cost function satisfying that $c_i(\top, e) = 0$ and that $(e, e') \in Q_i \cup Q_i^*$ implies $c_i(\varphi, e) = c_i(\varphi, e')$, for all $i \in I$, $\varphi \in \mathcal{L}$.

A question map $\mathcal{C} : I \rightarrow \mathcal{L}$ assigns each agent a formula.

For $e \in E$, a triple $(\mathcal{A}, \mathcal{C}, e)$ is an attention action, with e the actual event. Let ALL be the class of attention actions.

The interpretation of attention actions is largely on par with the interpretation of epistemic actions from standard DEL. Preconditions state the conditions under which events can occur, and the Q_i relation represents *unavoidable* indistinguishability: If $(e, e') \in Q_i$, then i simply cannot distinguish e and e' .




Fig. 1. Example for a single agent i . *Left*: attention models \mathcal{A} and \mathcal{A}' . Events are outlined and labeled with their preconditions. Actual events are thickly outlined. Dotted lines show Q_i^* , with reflexive links omitted (Q_i has only reflexive links). Costs functions c_i, c'_i are assumed constant across events. Assume $\forall \varphi, \psi \in \mathcal{L}, \forall e \in \mathcal{A}, \mathcal{A}', c_i(\varphi + \psi, e) = c_i(\varphi, e) + c_i(\psi, e)$, and the same for c' . *Top Right* and *Bottom Right*: attention state (M, w) and updates. Worlds are labeled with the literals they satisfy, the actual world is underlined. Full lines show R_i . The attention resource A_i is constant across worlds. Arrows show updates and are labeled with the action model and question used. **Story:** Agent i is uncertain of the truth value of theorems p and q , but has the goal to learn them. She can either do proofs (\mathcal{A}), or look up known implications (\mathcal{A}'). Given her attention, can she succeed? *Top Right*: She attempts to prove $(p \wedge q)$, but drains her attention and fails to learn anything. *Bottom Right*: She successfully proves p , then successfully looks up $p \rightarrow q$. By this less attention-taxing strategy, she also learns q and so achieves her goal.

However, the interpretation of Q_i^* is novel: If $(e, e') \in Q_i^*$, then the events are by *default* indistinguishable for i , but may be distinguishable if i pays attention to the right question (and has the attention to do so). Through the product defined below, Q_i^* interacts with i 's question \mathfrak{C}_i , its cost, and i 's attention span: If i asks a question which e and e' answer differently (i.e. $pre(e) \models \mathfrak{C}_i$ and $pre(e') \not\models \mathfrak{C}_i$) and for which i has sufficient attention ($A_i(w) \geq c_i(\mathfrak{C}_i, e)$), then e and e' will be distinguishable for i in world w . This reduces i 's attention in the updated model. Hence agents may learn the answers to their questions by paying attention.

About costs, we only assume that agents know the cost of each formula, and that asking no question (\top) is free. One could also require e.g. that $c_i(K_j \varphi, e) \geq c_i(\varphi, e)$ or $c_i(\varphi \wedge \psi, e) \geq c_i(\varphi, e) + c_i(\psi, e)$. Such assumptions would not influence the paper's results. About questions, we make the strong assumption that what is payed attention to is common knowledge. In Sec. 5, we remark on lifting this assumption.

Definition 8. Let $(M, w) = ((W, R, V, A), w)$ be an attention state and let $X = (\mathcal{A}, \mathfrak{C}, e)$ be an attention action with $\mathcal{A} = (E, Q, Q^*, pre, c)$. For all $i \in I$ and $\varphi \in \mathcal{L}$, let

$$Q_i^*[\varphi | \neg \varphi] = Q_i^* \setminus \{(e, e') \in Q_i^* : pre(e) \models \varphi \text{ and } pre(e') \not\models \varphi\}.$$

The attention update of (M, w) with X is $(M, w) \otimes X = ((W^X, R^X, V^X, A^X), (w, e))$ where

$$\begin{aligned} W^X &= \{(w, e) \in W \times E : M, w \models \text{pre}(e)\}, \\ R^X &\text{ is given by } (w, e)R_i^X(v, f) \text{ iff } (w, v) \in R_i \text{ and} \\ (e, f) &\in \begin{cases} Q_i \cup Q_i^* & \text{if } c_i(\mathfrak{E}_i, e) > A_i(w) \\ Q_i \cup Q_i^* / \mathfrak{E}_i \mid \neg \mathfrak{E}_i & \text{if } c_i(\mathfrak{E}_i, e) \leq A_i(w) \neq 0 \end{cases} \\ V^X(p) &= \{(w, e) \in W^X : w \in V(p)\} \text{ for all } p \in \Phi, \\ A_i^X(w, e) &= \max\{0, A_i(w) - c_i(\mathfrak{E}_i, e)\}, \end{aligned}$$

Call X applicable to (M, w) if $M, w \models \text{pre}(e)$, else not. Where σ is a (potentially infinite) sequence of attention actions, if for all $k \leq n$, X_{k+1} is applicable to $(M, w) \otimes X_1 \otimes \dots \otimes X_k$, denote $(M, w) \otimes X_1 \otimes \dots \otimes X_n$ by $(M, w)^\sigma = ((W^\sigma, R^\sigma, V^\sigma, A^\sigma), w^\sigma)$ and call σ applicable to (M, w) .

2.4 No Free Lunch

The attention actions of Figures 1 and 2 belong to a class shown special interest in this paper: the class of *No Free Lunch* (NFL) actions.

Definition 9. Let NFL be the class of attention actions $(\mathcal{A}, \mathfrak{E}, e) = ((E, Q, Q^*, \text{pre}), \mathfrak{E}, e)$ that satisfy, for all $i \in I$, $Q_i^* = E \times E$ and for all $e \in E$, $\varphi \in \mathcal{L} \setminus \{\top\}$, $c_i(\varphi, e) > 0$.

We find NFL actions to be of special interest as they enforce attention use for all non-trivial questions, and thus respect that attention is a bounded resource for learning. Jointly, the two restrictions entail that any change in information comes at some cost to attention, with the exception that agents always learn an NFL action's unavoidable "background information" $\bigvee_{e \in E} \text{pre}(e)$. Special cases where $\top \models \bigvee_{e \in E} \text{pre}(e)$ then ensure no learning without attention cost, as in $\mathcal{A}, \mathcal{A}'$ of Figure 1 and $\mathcal{A}, \mathcal{A}''$ of Figure 2. Stated differently, then an NFL action applied without attention cost is equivalent to the *public announcement* of its background information. This is the content of Lemma 11 below, used later to show decidability.

Definition 10. For any $X = ((E, Q, Q^*, \text{pre}, c), \mathfrak{E}, e)$, its background announcement is $X^! = ((E^!, Q^!, Q^{*!}, \text{pre}^!, c^!), \mathfrak{E}^!, e^!)$ for $E^! = \{e^!\}$, $Q_i^! = Q_i^{*!} = \{(e^!, e^!)\}$, $\text{pre}^!(e^!) = \bigvee_{e \in E} \text{pre}(e)$, $c^! = c$ and $\mathfrak{E}_i^! = \top$, for all $i \in I$.

Lemma 11. For any $X \in \text{NFL}$, for any attention state $(M, w) = ((W, R, V, A), w)$, if $A^X = A$, then $(M, w)^X \Leftrightarrow (M, w)^{X^!}$.

Proof. Let $X = ((E, Q, Q^*, \text{pre}, c), \mathfrak{E}, e)$. Then $Z = \{(w, e), (w, e^!)\} : w \in W, e \in E\} \subseteq (M, w)^X \times (M, w)^{X^!}$ is a bisimulation: **Atoms:** for all $p \in \Phi$, $w \in W$, $e \in E$, $i \in I$, $(w, e) \in V^X(p)$ iff $(w, e^!) \in V^{X^!}(p)$ and $A_i^X(w, e) = A_i^{X^!}(w, e^!)$ as both $A^X = A$ and $A^{X^!} = A$, the latter as $\mathfrak{E}_i^! = \top$ for all $i \in I$. **Forth:** Assume $(w, e)Z(w, e^!)$ and $(w, e)R_i^X(v, f)$ for some $i \in I$. Then wR_iv and as $Q_j^! = Q_j^{*!} = \{(e^!, e^!)\}$ for all $j \in I$, so $(w, e^!)R_i^{X^!}(v, e^!)$. Hence $(v, f)Z(v, e^!)$.

Back: Assume $(w, e)Z(w, e^!)$ and $(w, e^!)R_i^{X^!}(v, e^!)$ for some $i \in I$. Then wR_iv . As $(v, e^!) \in (M, w)^{X^!}$, $M, v \models \text{pre}^!(e^!)$, i.e. $M, v \models \bigvee_{e \in E} \text{pre}(e)$ (by Defs. 10 and 8). Hence, for some $f \in E$, $M, v \models \text{pre}(f)$. So $(v, f) \in (M, w)^X$. As $X \in \text{NFL}$, eQ^*f . Hence, also $(w, e)R_i^X(v, f)$ (cf. Def. 8) as $A^X = A$ implies that either $\mathfrak{C}_i = \top$ or $A_i(w) = 0 < c_i(\mathfrak{C}_i, e)$. Finally, $(v, f)Z(v, e^!)$. Hence, $(M, w)^X \Leftrightarrow (M, w)^{X^!}$.

Requirement $A^X = A$ of Lemma 11 states that no agent spends attention. We remark that if agent i spends attention and learns, then, even if j does not spend attention, j 's higher-order will still change. This is as in standard DEL where any informational change one way or the other affects all agents.




Fig. 2. Three variants of Muddy Children for $I = \{a, b, c\}$. Let φ_C be short for $\bigwedge_{i \in I} \neg K_i d \wedge \neg K_i \neg d$: No child knows whether they are dirty (d). *On the left column:* Three attention action models. \mathcal{A} is a public announcement of φ_C , the update used in the standard modeling of Muddy Children. In \mathcal{A}' , the children may fail to pay attention to the announcement. \mathcal{A}'' is as \mathcal{A}' , except that c cannot hear the announcement (the full line is the Q -relation). *Rows 1, 2, and 3* show different dynamics, each starting *after* the parent's announcement that at least one child is dirty. Each world w is labeled $A_a(w)A_b(w)A_c(w)$, with $A_i(w)$ dotted if i is dirty in w . Horizontal lines show R_a , vertical lines R_b , and diagonal lines R_c . The arrows show updates and are labeled with the action model and the question used by all agents. *Row 1:* Standard Muddy Children (see e.g. [13]): The children pay only attention to \top , but are forced to learn. *Row 2:* Muddy Children with limited attention: The children all pay attention to φ_C . After the first announcement, a runs out of attention, so does not learn its own state by the second announcement. *Row 3:* As Row 2, but a starts with zero attention and c cannot hear the announcement, despite paying attention. Nobody learns their state.

3 Emulation Results

This section shows that attention states may be recast as DEL epistemic states and *vice versa*, that every DEL epistemic action *without* postconditions may be emulated by an attention action but not *vice versa*, and that every attention action may be emulated by an epistemic action *with* postconditions, but not *vice versa*. These results facilitate proving the above modal equivalence result (Prop. 4) and undecidability of general attention planning (Thm. 24) as corollaries to existing results. They further show the relationship with DEL and allow us to place our results on epistemic planning with attention in the wider context of DEL-based epistemic planning in Sec. 5.

3.1 Attention States and Epistemic States

Definition 12. An epistemic Kripke model for I and At is a tuple $K = (W, R, V)$ where $W \neq \emptyset$ is a finite set of worlds, $R : I \rightarrow \mathcal{P}(W^2)$ assigns each $i \in I$ an equivalence relation R_i ; $V : At \rightarrow \mathcal{P}(W)$ is a valuation. For $w \in W$, (K, w) is called an epistemic state.

Truth of \mathcal{L} -formulas over (K, w) is as in Def. 2, except for the atomic clause:

$K, w \models p$ iff $w \in V(p)$ for all $p \in \Phi \cup \Psi$.

The Kripke rendition of the attention state $(M, w) = ((W, R, V, A), w)$ is $K(M, w) = ((W, R, V_A), w)$ where for all $p \in \Phi \cup \Psi$, $V_A(p) = \{w \in W : M, w \models p\}$.

Our first emulation result shows that the class of attention states can be embedded in the class of epistemic states:

Proposition 13. For any attention state (M, w) , $K(M, w)$ is *i)* an epistemic state; *ii)* isomorphic to (M, w) ; *iii)* for all $\varphi \in \mathcal{L}$, $K(M, w) \models \varphi$ iff $(M, w) \models \varphi$.

Proof. Immediate. See Proof Appendix for details.

Prop. 13 allows us to establish Prop. 4 as a corollary. For standard definitions of a bisimulation for Kripke models and the corresponding Hennessy-Milner Theorem, see [5, Def. 2.16 and Thm. 2.24].

Proof (Proof of Prop. 4). Assume that (M, w) and (M', w') are bisimilar witnessed by Z . Then Z is also a (Kripke model) bisimulation between $K(M, w)$ and $K(M', w')$, as is evident from their definition. By Prop. 13, if $(M, w) \models \varphi$, then also $K(M, w) \models \varphi$, and by the Hennessy-Milner Theorem, $K(M', w') \models \varphi$. Again by Prop. 13, $(M', w') \models \varphi$.

Our second emulation result inverts the first, showing that the class of epistemic states can be embedded in the class of attention states. As attention states have specific conditions for attention atom satisfaction, equi-satisfaction holds only for the sublanguage $\mathcal{L}_\Phi \subseteq \mathcal{L}$ of formulas without attention atoms. As Φ can be arbitrarily extended, we ignore this point when we use the result in showing undecidability (Thm. 24).

Proposition 14. *For every epistemic state (K, w) , there is an attention state (M, v) such that for all $\varphi \in \mathcal{L}_\Phi$, $(K, w) \models \varphi$ iff $(M, v) \models \varphi$.*

Proof. Adding to K any attention resource function A yields an attention model M for which (M, w) is as desired.

3.2 Attention Actions and Epistemic Actions

Definition 15. *An action model is a tuple $\mathcal{E} = (E, Q, pre, post)$ with $E \neq \emptyset$ a finite set of events; $Q: I \rightarrow \mathcal{P}(E^2)$ assigns each $i \in I$ an equivalence relation Q_i ; $pre: E \rightarrow \mathcal{L}$, $post: E \rightarrow (At \rightarrow \mathcal{L})$ assign each action a pre- and a post-condition, respectively. \mathcal{E} is without postconditions if for all $e \in E$, $post(e)(p) = p$.*

For $e \in E$, (\mathcal{E}, e) is an epistemic action, with e the actual event. Let POST and NOPOST denote, respectively, the classes of all epistemic actions all epistemic actions without postconditions.

Definition 16. *Let $(K, w) = ((W, R, V), w)$ be an epistemic state for At and $X = ((E, Q, pre, post), e)$ an epistemic action. The product update of (K, w) with $X = (\mathcal{E}, e)$ is $(K, w)^X = ((W^X, R^X, V^X), (w, e))$, where $W^X = \{(w, e) \in W \times E : M, w \models pre(e)\}$, $((w, e), (v, f)) \in R_i^X$ iff $(w, v) \in R_i$ and $(e, f) \in Q$ for all $i \in I$, and $V^X(p) = \{(w, e) \in W^X : M, w \models post(e)(p)\}$. Definitions of, and notation for, applicability is parallel to Def. 8.*

Definition 17. *For $Y = (\mathcal{A}, \mathcal{E}, e') \in \text{ALL}$, $X = (\mathcal{E}, e) \in \text{POST} \cup \text{NOPOST}$, Y is equivalent to X iff for all attention states (M, w) , all $\varphi \in \mathcal{L}$, $(M, w)^Y \models \varphi$ iff $(K(M, w))^X \models \varphi$.*

Proposition 18. *For all $X \in \text{NOPOST}$, there is an equivalent $Y \in \text{ALL}$.*

Proof (Proof sketch). Let $X = ((E, Q, pre, post), e) \in \text{NOPOST}$. To emulate X , we build an attention action $Y = ((E', Q', Q^*, pre', c), \mathcal{E}, e')$ where all agents pay attention to \top , and $Q^* = \{(e, e) : e \in E\}$. Then Y is essentially reduced to an epistemic action from NOPOST. If $E' = E$, $Q' = Q$, $pre' = pre$, $e' = e$, then Y is equivalent to X .

Jointly, Propositions 13 and 18 show that attention states and actions can emulate the dynamics of DEL, for epistemic actions in NOPOST. The converse of Prop. 18 is instead false: The dynamics of standard DEL cannot emulate those of attention states and actions, if one restricts to epistemic actions in NOPOST.

Proposition 19. *There is a $X \in \text{ALL}$ with no equivalent $Y \in \text{NOPOST}$.*

Proof (Proof sketch). Attention actions may change the valuation of attention atoms, which epistemic actions in NOPOST cannot.

If we do not restrict to NOPOST, then DEL can emulate all dynamics induced by attention actions:

Proposition 20. *For all $X \in \text{ALL}$, there is an equivalent $Y \in \text{POST}$.*

Proof (Proof sketch). Let $X = ((E, Q, Q^*, pre, c), \mathfrak{E}, e)$. We build an exponentially larger $Y = ((E', Q', pre', post), e) \in \text{POST}$. Let 2^I be the set of maps $In : I \rightarrow \{0, 1\}$. Intuitively, $In(i) = 1$ represents that i has enough attention to pay for \mathfrak{E}_i . Let $E' = \{f_{In} : f \in E, In \in 2^I\}$. Setting $pre'(f_{In})$ to

$$pre(f) \wedge \bigwedge_{In(i)=1} \alpha_i \geq c_i(\mathfrak{E}_i, f) \wedge \bigwedge_{In(j)=0} \alpha_j < c_j(\mathfrak{E}_j, f)$$

ensures that E' contains for each $f \in E$ and each $In \in 2^I$ a unique event f_{In} with $pre'(f_{In})$ satisfied at worlds that satisfy $pre(f)$ and where exactly agents in $\{i \in I : In(i) = 1\}$ have attention exceeding the cost of \mathfrak{E}_i . The construction ensures that for any attention state (M, w) , $(w, f_{In}) \in K(M, w)^Y$ iff $(w, f) \in (M, w)^X$ and $In(i) = 1$ iff $A_i(w) \geq c_i(\mathfrak{E}_i, f)$. (M, w) and $K(M, w)$ thus have the same number of worlds, and for all $p \in \mathcal{P}$, (w, f_{In}) satisfies p iff (w, f) does.

To make Y correctly update the attention atoms, we use that \mathcal{L} can express “ n will be the attention value of agent i after the event f ”, writing $next_f(\alpha_i = n)$ as abbreviation for

$$\begin{aligned} & ((\alpha_i = n) \wedge (\alpha_i = \max\{0, n - c_i(\mathfrak{E}_i, f)\})) \vee \\ & (\neg(\alpha_i = n) \wedge (\alpha_i = \min\{n + c_i(\mathfrak{E}_i, f), N\})) \end{aligned}$$

The first disjunct reads “ n is i ’s current attention value and $c_i(\mathfrak{E}_i, f)$ —the cost of what i pays attention to in f —is zero” and the second “ n is not the current attention value, but n is the current value minus $c_i(\mathfrak{E}_i, f)$ ”. We use $next_f(\alpha_i = n)$ in assigning postconditions, with two key clauses being $post(f_{In})(\alpha_i = n) = next_f(\alpha_i = n)$ and $post(f_{In})(\alpha_i < n) = \bigvee_{0 \leq j < n} next_f(\alpha_i = j)$, for all $f_{In} \in E'$.

Finally, we let $(f_{In}, g_{In'}) \in Q'$ iff either (1) $(f, g) \in Q_i$ or (2) $(f, g) \in Q_i^*$ and $pre'(f_{In}) \models \alpha_j < c_i(\mathfrak{E}_i, f)$, or (3) $(f, g) \in Q_i^*$ and $pre'(f_{In}) \models \mathfrak{E}_i$ iff $pre'(g_{In'}) \models \mathfrak{E}_i$. This ensures that (w, f) and (v, g) are related for i in $(M, w)^X$ iff (w, f_{In}) and $(v, g_{In'})$ are related for i in $K(M, w)^Y$.

The construction ensures that $(M, w)^X$ and $K(M, w)^Y$ are isomorphic, entailing that X and Y are equivalent.

To finalize our emulation results, we show Prop. 20’s converse false.

Proposition 21. *There exists a $X \in \text{POST}$ with no equivalent $Y \in \text{ALL}$.*

Proof (Proof sketch). Epistemic actions may change the valuation of atoms other than the attention atoms. This cannot be done by attention actions.

4 Epistemic Planning with Attention

Finally, we turn to epistemic planning with attention. The following definitions follow those for DEL-based epistemic planning, as in [8].

Definition 22. *An attention planning task $T = (s_0, \Sigma, \varphi_g)$ consists of an (initial) attention state s_0 ; a finite set of attention actions Σ ; and a goal formula $\varphi_g \in \mathcal{L}$. A solution to T is a finite sequence X_1, \dots, X_n of actions from Σ applicable to s_0 such that $s_0^{X_1, \dots, X_n} \models \varphi_g$.*

Definition 23. For $X \subseteq \text{ALL}$, denote by \mathcal{T}_X the class of attention planning tasks $T = (s_0, \Sigma, \varphi_g)$ with $\Sigma \subseteq X$. Given $\mathcal{T}_X \subseteq \mathcal{T}_{\text{ALL}}$, denote by $\text{PlanEx-}\mathcal{T}_X$ the decision problem (called the plan existence problem on \mathcal{T}_X): Given a planning task $T \in \mathcal{T}_X$, does T have a solution?

For DEL-based epistemic planning, the general plan existence problem is undecidable [7]. The same holds when attention is involved:

Theorem 24. $\text{PlanEx-}\mathcal{T}_{\text{ALL}}$ is undecidable.

Proof. The plan existence problem for DEL epistemic actions with preconditions of modal depth at most n and without postconditions $\text{PlanEx-}\mathcal{T}(n, -1)$ is undecidable for $n \geq 2$ [8]. By Propositions 14 and 18, for each epistemic planning task in $\text{PlanEx-}\mathcal{T}(n, -1)$, there is an equivalent attention planning task in $\text{PlanEx-}\mathcal{T}_{\text{ALL}}$, which is hence undecidable.

For the class of NFL attention actions, we obtain a more encouraging result:

Theorem 25. $\text{PlanEx-}\mathcal{T}_{\text{NFL}}$ is decidable.

Proof. Let $T = (s_0, \Sigma, \varphi_g)$ be an attention planning task with $\Sigma \subseteq \text{NFL}$. We show that checking if T has a solution is decidable by showing the claim that for any sequence $(X_k)_{k \in \mathbb{N}}$ of Σ -actions applicable to s_0 , there is an $n \in \mathbb{N}$ such that for all $Y \in \Sigma$, $s_0^{X_1, \dots, X_n} \not\leftrightarrow s_0^{X_1, \dots, X_n, Y}$. Hence, only a finite set of plans needs to be checked to determine if a solution exists.

We tacitly identify attention states with their bisimulation contractions (Def. 5), justified by Lemma 6. This makes the reference to the cardinality of sets of worlds meaningful.

Let $(M, w) = (W, R, V, A, w)$ be any attention state. Applying any $X \in \Sigma$ to (M, w) will either consume attention for some agent (so $A^X \neq A$), or it will not consume any attention ($A = A^X$). The claim then follows from two points:

First, as attention is finite, any applicable sequence of actions will eventually stop consuming it: For any $X \in \Sigma$, if $A^X \neq A$, then for some $i \in I$, some $w \in W$ and some event e from X , $A_i^X(w, e) < A_i(w)$. Hence, as $\{A_i(v) : i \in I, v \in W\}$ is bounded, for any sequence $(X_k)_{k \in \mathbb{N}}$ of Σ -actions applicable to (M, w) , there is an $n \in \mathbb{N}$ such that for all $X \in \Sigma$, $A^{X_1, \dots, X_n} = A^{X_1, \dots, X_n, X}$.

Second, as W is finite, any sequence of actions none of which consume attention will eventually reach a fixed point: Let $(X_k)_{k \in \mathbb{N}}$ be a sequence of Σ -actions applicable to (M, w) such that $A^{X_1, \dots, X_k} = A^{X_1, \dots, X_k, X_{k+1}}$. By Lemma 11, $|W^{X_1, \dots, X_k, X_{k+1}}| < |W^{X_1, \dots, X_k}|$ (as we identify models with their bisimulation contraction and as each $X_k!$ has a single event $e!$) or $(M, w)^{X_1, \dots, X_k, X_{k+1}} \not\leftrightarrow (M, w)^{X_1, \dots, X_k}$ (if $M, v \models \text{pre}(e!)$ for all $v \in W$). As W is finite, the first disjunct can eventually not obtain, so a fixed point is reached.

The two points jointly imply the claim: for any sequence $(X_k)_{k \in \mathbb{N}}$ of Σ -actions applicable to s_0 , the first ensures that after some $n \geq 0$ steps, no more attention is consumed, and the second implies that after (additional) $m \geq 0$ steps, a fixed point is reached. Hence, finally, if φ_g is not reached by the fixed point, $(X_j)_{j \leq n+m}$ is not a solution.

5 Final Remarks

Our main results show that the plan existence problem for epistemic planning with attention is in general undecidable (Theorem 24), but decidable for No Free Lunch (NFL) actions (Theorem 25). As NFL actions are of special interest, this is an encouraging result for epistemic planning with attention.

As (DEL-based) epistemic planning is in general undecidable, the emulation results of Sec. 3 makes Theorem 24 unsurprising. However, Theorem 25 strikes a sharp contrast with other results in epistemic planning:

Epistemic planning task classes have earlier been investigated by number of agents, with the single-agent case decidable [7], and the n -agent $n \geq 2$ cases undecidable [2] (both without postconditions). Ensuing, decidable fragments have been sought in the hierarchy of classes $\mathcal{T}(m, n)$ allowing epistemic actions with modal depth of pre- and postconditions at most m and n , respectively. [8] provide a recent survey, contributing so that the only open question concerns $\mathcal{T}(1, -1)$, ‘-1’ referring to no postconditions. Of special interest to this paper are the classes that attention planning can emulate ($\mathcal{T}(m, 0)$) or can be emulated by ($\mathcal{T}(m, 1)$). The case for $\mathcal{T}(0, 0)$ is decidable, but undecidable for $\mathcal{T}(m, 0), m \geq 1$ and $\mathcal{T}(m, 1), m \geq 0$ [8]. However, the emulations of Sec. 3 makes it evident that the class $\mathcal{T}_{\text{NFL}}^{\text{DEL}}$ of epistemic planning tasks obtained by emulating \mathcal{T}_{NFL} is a proper subclass of $\bigcup_{m \in \mathbb{N}} \mathcal{T}(m, 0) \cup \mathcal{T}(m, 1)$ with $\mathcal{T}_{\text{NFL}}^{\text{DEL}} \cap (\mathcal{T}(m, 0) \cup \mathcal{T}(m, 1))$ non-empty for every $m \in \mathbb{N}$. By Theorem 25, the plan existence problem for $\mathcal{T}_{\text{NFL}}^{\text{DEL}}$ is decidable, for any number of agents. Hence, the paper’s results finds a class of epistemic planning tasks that cuts across the oft studied classes, for which the plan existence problem is decidable, and which, by allowing arbitrary preconditions, is still reasonably expressive.

We conclude with remarks on future research. First, regarding complexity, then we venture that the plan existence problem for \mathcal{T}_{NFL} is NP-complete, as this is the case for epistemic planning with public announcements [10]. Second, we believe \mathcal{L} is sufficiently expressive to obtain a complete axiom system for attention states. This, the emulations and existing completeness results for epistemic actions with postconditions [12] would yield a complete dynamic logic for attention actions. However, the emulations enforce an exponential blowup that would spill into reduction axioms. It may therefore be desirable to define such directly for attention actions. Finally: The model assumes common knowledge of what agents pay attention to. This may be dropped by adding indistinguishability between attention maps, at the cost of a more elaborate product. We have left this construction for longer work.

Acknowledgments. We thank the Center for Information and Bubble Studies, funded by the Carlsberg Foundation. RKR was partially supported by the DFG-ANR joint project *Collective Attitude Formation* [RO 4548/8-1].

References

1. Alechina, N., Logan, B., Whitsey, M.: A Complete and Decidable Logic for Resource-Bounded Agents. In: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems. pp. 606–613 (2004)
2. Aucher, G., Bolander, T.: Undecidability in Epistemic Planning. In: Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI-13. pp. 27–33 (2013)
3. Balbiani, P., Fernández-Duque, D., Lorini, E.: The Dynamics of Epistemic Attitudes in Resource-Bounded Agents. *Studia Logica* **107**, 457–488 (2019)
4. Baltag, A., Moss, L.S., Solecki, S.: The Logic of Public Announcements, Common Knowledge, and Private Suspicions (extended abstract). In: Proceedings of the Seventh Conference on Theoretical Aspects of Rationality and Knowledge, TARK-98. pp. 43–56 (1998)
5. Blackburn, P., de Rijke, M., Venema, Y.: *Modal Logic*. Cambridge University Press (2001)
6. Bolander, T.: A gentle introduction to epistemic planning: The DEL approach. In: *Electronic Proceedings in Theoretical Computer Science*. vol. 243, pp. 1–22 (2017)
7. Bolander, T., Birkegaard, M.: Epistemic planning for single- and multi-agent systems. *Journal of Applied Non-Classical Logics* **21**(1), 9–34 (2011)
8. Bolander, T., Charrier, T., Pinchinat, S., Schwarzenrüber, F.: DEL-based epistemic planning: Decidability and complexity. *Artificial Intelligence* **287**, 103304 (2020)
9. Bolander, T., van Ditmarsch, H., Herzig, A., Lorini, E., Pardo, P., Schwarzenrüber, F.: Announcements to Attentive Agents. *Journal of Logic, Language and Information* **25**, 1–35 (2016)
10. Bolander, T., Jensen, M.H., Schwarzenrüber, F.: Complexity Results in Epistemic Planning. In: Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-15 (2015)
11. van Ditmarsch, H., Herzig, A., Lorini, E., Schwarzenrüber, F.: Listen to Me! Public Announcements to Agents that Pay Attention - or Not. In: *International Workshop on Logic, Rationality and Interaction*. pp. 96–109 (2013)
12. van Ditmarsch, H., Kooi, B.: Semantic Results for Ontic and Epistemic Change. In: *Logic and the Foundations of Game and Decision Theory*. pp. 87–117 (2008)
13. Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.Y.: *Reasoning About Knowledge*. The MIT Press (1995)
14. Hefti, A., Heinke, S.: On the economics of superabundant information and scarce attention. *Economia. History, Methodology, Philosophy* (5-1), 37–76 (2015)
15. Lorini, E., Tummolini, L., Herzig, A.: Establishing Mutual Beliefs by Joint Attention: Towards a Formal Model of Public Events. In: *Proceedings of the Annual Meeting of the Cognitive Science Society*. pp. 1325–1330 (2005)
16. Rendsvig, R., Symons, J.: Epistemic Logic. In: Zalta, E.N. (ed.) *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2021 edn. (2021)
17. Schipper, B.C.: Awareness. In: van Ditmarsch, H., Halpern, J.Y., van der Hoek, W., Kooi, B.P. (eds.) *Handbook of Epistemic Logic* (2014)
18. Simon, H.A.: Designing organizations for an information-rich world. In: Greenberger (ed.) *Computers, communications, and the public interest*, pp. 37–72 (1971)
19. Smets, S., Solaki, A.: The Effort of Reasoning: Modelling the Inference Steps of Boundedly Rational Agents. In: Moss, L.S., de Queiroz, R., Martinez, M. (eds.) *Logic, Language, Information, and Computation*. pp. 307–324 (2018)