Towards Epistemic-Doxastic Planning with Observation and Revision

Thorsten Engesser[1], Andreas Herzig[1], Elise Perrotin[2]

KR and MAS Conventicle 2024

[1]IRIT, CNRS, Toulouse, France
[2]CRIL, CNRS, Lens, France

Motivation

Planning with a theory of mind is a valuable skill for autonomous agents:

- Accounting for other agents with *false beliefs*.
- Planning to facilitate *coordination*.

Most existing planning formalisms support *knowledge* or *belief*, but not both.

Our approach is inspired by *lightweight* epistemic & doxastic planning approaches from the literature [Cooper et al., 2021, Muise et al., 2022].

Epistemic-Doxastic Logic S5-EDL

We consider the epistemic-doxastic logic S5-EDL:

$$\varphi ::= K_i\varphi \mid B_i\varphi \mid \neg\varphi \mid \varphi \wedge \varphi$$

- Knowledge: Facts which we currently observe.
- Belief: Things we observed in the past or learn through communication.
- S5 for knowledge, KD45 for belief + interactions axioms (e.g., $K_i\varphi \rightarrow B_i\varphi$).
- The satisfiability problem of the full logic is PSPACE-complete.

We do not want to use Kripke models + DEL update models for states/actions.

$\Rightarrow$ Can we find something simpler?

Knowledge Only: The Epistemic Logic of Observations (EL-O)

States = Valuations over observation atoms [Cooper et al., 2021]:

$$\sigma ::= p \mid S_i\sigma$$

- $S_i\sigma$: *agent i sees* σ (= *knowing whether*, $S_i\sigma \equiv K_i\sigma \vee K_i\neg\sigma$).
- No negations, conjunctions or disjunctions within modal operator.
- Introspective atoms such as S_1S_1p are excluded (they are tautological).

Example:

$$\{p, S_1p, S_1S_2p\} \models K_1p \wedge \neg K_2p \wedge K_1\neg K_2p$$

Knowledge Only: The Epistemic Logic of Observations (EL-O)

- Introspection-free observation atoms logically independent of each other.

$\Rightarrow$ S5-satisfiability of formulas over such atoms reduces to boolean SAT.

Unfortunately, this approach does not work with *having a belief about*:

$$B_i\varphi \not\equiv BA_i\varphi \wedge \varphi$$

4 epistemic situations:

	$\neg\sigma$	σ
$\neg S_i\sigma$	$\neg\sigma \wedge \neg K_i\neg\sigma$	$\sigma \wedge \neg K_i\sigma$
$S_i\sigma$	$\neg\sigma \wedge K_i\neg\sigma$	$\sigma \wedge K_i\sigma$

6 doxastic situations:

	$\neg\beta$	β
?	$\neg\beta \wedge \neg B_i\neg\beta \wedge \neg B_i\beta$	$\beta \wedge \neg B_i\neg\beta \wedge \neg B_i\beta$
?	$\neg\beta \wedge \neg B_i\neg\beta \wedge B_i\beta$	$\beta \wedge \neg B_i\neg\beta \wedge B_i\beta$
?	$\neg\beta \wedge B_i\neg\beta \wedge \neg B_i\beta$	$\beta \wedge B_i\neg\beta \wedge \neg B_i\beta$

True Beliefs and Mere Beliefs [Herzig and Perrotin, 2021]

True belief about φ: $\mathsf{TBA}_i\varphi \equiv (B_i\varphi \wedge \varphi) \vee (B_i\neg\varphi \wedge \neg\varphi)$

Mere belief about φ: $\mathsf{MBA}_i\varphi \equiv (B_i\varphi \wedge \neg K_i\varphi) \vee (B_i\neg\varphi \wedge \neg K_i\neg\varphi)$

All combinations of knowledge and belief are expressible:

		e.g., assuming φ is true:
i has no belief about φ	$\neg\mathsf{MBA}_i\varphi \wedge \neg\mathsf{TBA}_i\varphi$	$\varphi \wedge \neg B_i\varphi \wedge \neg B_i\neg\varphi$
i has a false belief about φ	$\mathsf{MBA}_i\varphi \wedge \neg\mathsf{TBA}_i\varphi$	$\varphi \wedge B_i\neg\varphi \wedge \neg K_i\varphi$
i has a lucky belief about φ	$\mathsf{MBA}_i\varphi \wedge \mathsf{TBA}_i\varphi$	$\varphi \wedge B_i\varphi \wedge \neg K_i\varphi$
i knows whether / observes φ	$\neg\mathsf{MBA}_i\varphi \wedge \mathsf{TBA}_i\varphi$	$\varphi \wedge B_i\varphi \wedge K_i\varphi$

A Lightweight Fragment of S5-EDL

We consider boolean formulas over so-called *REDA atoms*:

$$\alpha ::= p \mid \mathtt{TBA}_i\alpha \mid \mathtt{MBA}_i\alpha$$

- *REDA*: *repetition-free* epistemic-doxastic atoms.
  $\Rightarrow$ No negations, conjunctions or disjunctions within modal operator.
  $\Rightarrow$ Introspective atoms such as $\mathtt{TBA}_i\mathtt{MBA}_i\alpha$ are excluded.
- Arbitrary conjunctions of such atoms are satisfiable.
  $\Rightarrow$ Satisfiability reduces to propositional SAT (NP-complete).

We use valuations over *REDA* atoms as states. For example:

$$\{p, \mathtt{TBA}_i p, \neg\mathtt{MBA}_i p, \mathtt{TBA}_j p, \mathtt{MBA}_j p\} \models K_i p \wedge B_j p$$

Actions

Actions have indirect effects conditional on agents' observations.

E.g., action of changing the truth value of p:

- Direct effect: $\top \triangleright \pm p$
- Indirect effect: $\mathtt{MBA}_i p \triangleright \pm \mathrm{TBA}_i p$.
- Lucky beliefs become false beliefs and vice versa.
- There are additional higher-order indirect effects...

Actions

In our paper, we define the following types of actions:

- Ontic actions (changing the value of a proposition).
- Starting and stopping to observe (first and second-order).

Allows us to model some first- and second-order *false-belief tasks*.

Example: Sally-Anne Task

1. Two children, Sally and Anne, are in a room together.
2. Sally has a marble, which she puts into a basket.
3. Sally leaves the room to go out for a walk.
4. Anne removes the marble from the basket and puts it into a box.
5. Sally comes back into the room.

Will Sally search for her marble in the basket or in the box?

Example: Sally-Anne Task

$$\underset{\models \mathrm{OBS}_S p}{\boxed{\begin{array}{c} p, \\ \mathrm{TBA}_S p, \neg\mathrm{MBA}_S p, \\ \mathrm{TBA}_A p, \neg\mathrm{MBA}_A p \end{array}}} \xrightarrow[\text{observing}]{\text{Sally stops}} \underset{\models \mathrm{LBA}_S p}{\boxed{\begin{array}{c} p, \\ \mathrm{TBA}_S p, \underline{\mathrm{MBA}_S p}, \\ \mathrm{TBA}_A p, \neg\mathrm{MBA}_A p \end{array}}} \xrightarrow[\text{the marble}]{\text{Anne moves}} \underset{\models \mathrm{FBA}_S p}{\boxed{\begin{array}{c} \underline{\neg p}, \\ \underline{\neg\mathrm{TBA}_S p}, \mathrm{MBA}_S p, \\ \mathrm{TBA}_A p, \neg\mathrm{MBA}_A p \end{array}}}$$

If Sally starts to observe the marble again:

$$\xrightarrow[\text{observing}]{\text{Sally starts}} \underset{\models \mathrm{OBS}_S p}{\boxed{\begin{array}{c} \neg p, \\ \underline{\mathrm{TBA}_S p}, \underline{\neg\mathrm{MBA}_S p}, \\ \mathrm{TBA}_A p, \neg\mathrm{MBA}_A p \end{array}}}$$

We get *revision* for free!

Results (AAAI 2024)

- Satisfiability in our S5-EDL fragment reduces to propositional satisfiability.
- We define an epistemic-doxastic planning formalism.
- Planning reduces to classical planning (PSPACE-complete).

Limitations

We define only actions for second-order knowledge and beliefs.

- Could be generalized to higher-order.
- Actions with second-order indirect effects are already quite complicated.

Our approach only approximates second-order observability:

$$
\begin{aligned}
\mathsf{OBS}_i\mathsf{OBS}_j p &\equiv \mathsf{TBA}_i(\mathsf{TBA}_j p \wedge \neg\mathsf{MBA}_j p) \wedge \neg\mathsf{MBA}_i(\mathsf{TBA}_j p \wedge \neg\mathsf{MBA}_j p) \\
&\approx \mathsf{TBA}_i\mathsf{TBA}_j p \wedge \mathsf{TBA}_i\mathsf{MBA}_j p \wedge \neg\mathsf{MBA}_i\mathsf{TBA}_j p \wedge \neg\mathsf{MBA}_i\mathsf{MBA}_j p \\
&\equiv \mathsf{OBS}_i\mathsf{TBA}_j p \wedge \mathsf{OBS}_i\mathsf{MBA}_j p
\end{aligned}
$$

E.g., we cannot express:

"I know you don't observe p, but I have no idea what you believe."

Thank you!

References I

- Cooper, M. C., Herzig, A., Maffre, F., Maris, F., Perrotin, E., and Régnier, P. (2021).
  A lightweight epistemic logic and its application to planning.
  Artif. Intell., 298:103437.
- Herzig, A. and Perrotin, E. (2021).
  True belief and mere belief about a proposition and the classification of epistemic-doxastic situations.
  Filosofiska Notiser, 8(1):103–117.

References II

Muise, C., Belle, V., Felli, P., McIlraith, S. A., Miller, T., Pearce, A. R., and Sonenberg, L. (2022).
Efficient multi-agent epistemic planning: Teaching planners about nested belief.
Artif. Intell., 302:103605.