

Comprehensive Multi-Agent Epistemic Planning

Francesco Fabiano

Dipartimento DMIF, Università di Udine, Udine, Italy

francesco.fabiano@uniud.it

Over the last few years, the concept of Artificial Intelligence has become central in different tasks concerning both our daily life and several working scenarios. Among these tasks *automated planning* has always been central in the AI research community. In particular, this manuscript is focused on a specialized kind of planning known as *Multi-agent Epistemic Planning* (MEP).

Epistemic Planning (EP) refers to an automated planning setting where the agent reasons in the space of knowledge/beliefs states and tries to find a plan to reach a desirable state from a starting one. Its general form, the MEP problem, involves multiple agents who need to reason about both the state of the world and the information flows between agents. To tackle the MEP problem several tools have been developed and, while the diversity of approaches has led to a deeper understanding of the problem space, each proposed tool lacks some abilities and does not allow for a comprehensive investigation of the information flows. That is why, the objective of our work is to formalize an environment where a complete characterization of the agents' knowledge/beliefs interaction and update is possible. In particular, we aim to achieve such goal by defining a new action-based language for *multi-agent epistemic planning* and to implement an epistemic planner based on it. This solver should provide a tool flexible enough to reason on different domains, *e.g.*, economy, security, justice and politics, where considering others' knowledge/beliefs could lead to winning strategies.

1 Introduction

The proliferation of agent-based and IoT technologies has enabled the development of novel applications involving hundreds of agents. To operate in such domains agents need to make decisions based on both the available problem information and on their knowledge, or belief, about the preferences and actions of other agents involved in, or impacted by, their decisions.

To maximize the potentials of such autonomous systems, *multi-agent planning* and scheduling research [7–9, 23, 25] will need to keep pace. As already said, creating a plan for multiple agents to achieve a goal will need to take into consideration agents' knowledge, beliefs and to account for aspects like trust, dishonesty, deception and incomplete knowledge. The planning problem in this new setting is referred to as *Multi-agent Epistemic Planning* (MEP) in the literature. MEP is a generalization of the concept of *Epistemic Planning* (EP), which refers to a single agent setting that reasons about knowledge and beliefs. In the MEP scenario, agents reason about the information that other agents have of each other while also considering how such information is shared. That is, MEP is the research area that attempts to create artificial agents that build plans considering the information flows between the agents themselves.

Nevertheless, reasoning about knowledge and beliefs is not as direct as reasoning on the “physical” state of the world. In fact, expressing belief relations between a group of agents often implies to consider *nested* and *group* beliefs that are not easily extracted from the state description by a human reader. For this reason, it is necessary to develop a complete and accessible action language to model multi-agent epistemic domains [1] and to advance also in the study of epistemic solvers [1, 2, 12, 19, 22, 24, 29]. While these solvers may have different restrictions, *e.g.*, considering only the concept of *knowledge* while not

allowing the idea of *belief* or not allowing for *dynamic common knowledge/belief*, all of them share the same goal: to plan while considering the information flows between the agents.

In our research we are exploring the epistemic planning problem with particular focus on i) formalizing an updated version of the action languages $m\mathcal{A}^P$ [12] and E-PDDL [14]; and ii) on developing a comprehensive multi-agent epistemic solver flexible enough to be used in several scenarios such as: economy, security, justice and politics.

2 Background

This section mostly stems from [11] where the same concepts were introduced in order to describe the foundations of the work presented in this manuscript.

2.1 Dynamic Epistemic Logic

Logicians have always been interested in describing *the state of the world* through formalism that would allow to reason on the world with logic itself. This interest has led, among other things, to the formalization of the well-known planning problem [26] and to the introduction of several *modal logics* [5, 27, 28] used to describe different types of domains. *Dynamic Epistemic Logic* (DEL), in particular, is used to reason not only on the state of the world but also on *information change*. As said by Van Ditmarsch *et al.* [28]: “*information* is something that is relative to a subject who has a certain perspective on the world, called an *agent*, and that is meaningful as a whole, not just loose bits and pieces. This makes us call it *knowledge* and, to a lesser extent, *belief*”. The idea behind DEL is, therefore, to have a formalization that allows to reason on dynamic domains where, not only the state world is taken into consideration, but also the knowledge/beliefs that the agents have about the world and about the knowledge/beliefs of each other are considered.

2.1.1 Epistemic Logic

Dynamic Epistemic Logic is, clearly, connected to *epistemic logic*: that is the logic that allows to reason on the knowledge/belief of agents in static domains. This logic is based on two main concepts: i) *Kripke structures*, a data structure that is widely used in literature [1, 15, 28] to model its semantics; and ii) *belief formulae*, a type of formula that takes into consideration epistemic operators and is used to represent the knowledge/beliefs of the agents. As it is beyond the scope of this work to give an exhaustive introduction on epistemic logic let us provide only the fundamental definitions and intuitions; the reader who has interest in a more detailed description can refer to [15].

Let $\mathcal{A}\mathcal{G}$ be a set of agents and let \mathcal{F} be a set of propositional variables, called *fluents*. We have that each *world* is described by a subset of elements of \mathcal{F} (intuitively, those that are “true” in the world). Moreover, in epistemic logic each agent $ag \in \mathcal{A}\mathcal{G}$ is associated with an epistemic modal operator \mathbf{B}_{ag} that intuitively represents the knowledge/belief of ag . Finally, epistemic *group operators* \mathbf{E}_α and \mathbf{C}_α are also introduced in epistemic logic. Intuitively \mathbf{E}_α and \mathbf{C}_α represent the knowledge/belief of a group of agents α and the *common knowledge/belief* of α respectively. To be more precise, following the work of Baral *et al.* [1], we have that:

Definition 2.1. A *fluent formula* is a propositional formula built using the propositional variables in \mathcal{F} and the traditional propositional operators $\wedge, \vee, \Rightarrow, \neg$. We make use \top and \perp to indicate *True* and *False*, respectively. A *fluent atom* is a formula composed by just an element $f \in \mathcal{F}$, instead a *fluent literal* is

either a fluent atom $f \in \mathcal{F}$ or its negation $\neg f$. During this work let us refer to fluent literals simply as *fluents*.

Definition 2.2. A *belief formula* is defined as follow:

- A fluent formula is a belief formula;
- let φ be belief formula and $ag \in \mathcal{AG}$, then $\mathbf{B}_{ag}\varphi$ is a belief formula;
- let φ_1, φ_2 and φ_3 be belief formulae, then $\neg\varphi_3$ and $\varphi_1 \text{ op } \varphi_2$ are belief formulae, where $\text{op} \in \{\wedge, \vee, \Rightarrow\}$;
- all the formulae of the form $\mathbf{E}_\alpha\varphi$ or $\mathbf{C}_\alpha\varphi$ are belief formulae, where φ is itself a belief formula and $\emptyset \neq \alpha \subseteq \mathcal{AG}$.

From now on let us denote with $\mathcal{L}_{\mathcal{AG}}^{\mathbf{C}}$ the language of the belief formulae over the sets \mathcal{F} and \mathcal{AG} and with $\mathcal{L}_{\mathcal{AG}}$ as the language over beliefs formulae that does not allow the use of \mathbf{C} .

Example 2.1. Let us consider the formula $\mathbf{B}_{ag_1}\mathbf{B}_{ag_2}\varphi$. This formula expresses that the agent ag_1 believes that the agent ag_2 believes that φ is true, instead, $\mathbf{B}_{ag_1}\neg\varphi$ expresses that the agent ag_1 believes that φ is false.

As mentioned above the classical way of providing a semantics for the language of epistemic logic is in terms of *pointed Kripke structure* [21]. More formally:

Definition 2.3. A *Kripke structure* is a tuple $\langle S, \pi, \mathcal{B}_1, \dots, \mathcal{B}_n \rangle$, such that:

- S is a set of worlds;
- $\pi : S \mapsto 2^{\mathcal{F}}$ is a function that associates an interpretation of \mathcal{F} to each element of S ;
- for $1 \leq i \leq n$, $\mathcal{B}_i \subseteq S \times S$ is a binary relation over S .

Definition 2.4. A *pointed Kripke structure* is a pair (M, s) where M is a Kripke structure as defined above and $s \in S$, where s represents the real world.

Following the notation of [1], allow us to indicate with $M[S], M[\pi]$ and $M[i]$ the components S, π and \mathcal{B}_i of M , respectively. Intuitively $M[S]$ captures all the worlds that the agents believe to be possible and $M[i]$ encodes the beliefs of each agent. More formally the semantics on pointed Kripke structure is as follows:

Definition 2.5. Given the belief formulae $\varphi, \varphi_1, \varphi_2$, an agent ag_i , a group of agents α , a pointed Kripke structure $(M = \langle S, \pi, \mathcal{B}_1, \dots, \mathcal{B}_n \rangle, s)$:

- (i) $(M, s) \models \varphi$ if φ is a fluent formula and $\pi(s) \models \varphi$;
- (ii) $(M, s) \models \mathbf{B}_{ag_i}\varphi$ if for each t such that $(s, t) \in \mathcal{B}_i$ it holds that $(M, t) \models \varphi$;
- (iii) $(M, s) \models \mathbf{E}_\alpha\varphi$ if $(M, s) \models \mathbf{B}_{ag_i}\varphi$ for all $ag_i \in \alpha$;
- (iv) $(M, s) \models \mathbf{C}_\alpha\varphi$ if $(M, s) \models \mathbf{E}_\alpha^k\varphi$ for every $k \geq 0$, where $\mathbf{E}_\alpha^0\varphi = \varphi$ and $\mathbf{E}_\alpha^{k+1}\varphi = \mathbf{E}_\alpha(\mathbf{E}_\alpha^k\varphi)$;
- (v) the semantics of the traditional propositional operators is as usual.

2.1.2 Knowledge or Belief

As pointed out in the previous paragraph the modal operator \mathbf{B}_{ag} represents the worlds' relation in a Kripke structure and, as expected, different relations' properties imply different meaning for \mathbf{B}_{ag} . In particular in this work we are interested in representing the knowledge or the beliefs of the agents. The problem of formalizing these two concepts has been studied in depth bringing to an accepted formalization for both knowledge and beliefs. In particular, when a relation¹ respects all the axioms presented in Table 1 is called an **S5** relation and it encodes the concept of knowledge while when it encodes all the axioms but **T** it characterizes the concept of belief. Following this characterization we refer to knowledge and belief as **S5** and **KD45** logic respectively.

Property of \mathcal{B}	Axiom
$\mathcal{B}_i \varphi \Rightarrow \varphi$	T
$\mathcal{B}_i \varphi \Rightarrow \mathcal{B}_i \mathcal{B}_i \varphi$	4
$\neg \mathcal{B}_i \varphi \Rightarrow \mathcal{B}_i \neg \mathcal{B}_i \varphi$	5
$\neg \mathcal{B}_i \perp$	D
$(\mathcal{B}_i \varphi \wedge \mathcal{B}_i (\varphi \Rightarrow \psi)) \Rightarrow \mathcal{B}_i \psi$	K

Table 1: Knowledge and beliefs axioms. [15].

Intuitively the difference between the two logics is that an agent cannot *know* something that is not true in **S5** but she can *believe* it in **KD45**. As this introduction is not supposed to explore in depth this topic we will not go into further detail, and we address the interested reader to [15].

2.2 Multi-Agent Epistemic Planning

Epistemic planning [2] refers to the generation of plans for multiple agents to achieve goals which can refer to the state of the world, the beliefs of agents and/or the knowledge of agents. It has recently attracted the attention of researchers from various communities such as planning, dynamic epistemic logic and knowledge representation.

With the introduction of the classical planning problem, in the early days of artificial intelligence, several action languages (e.g., \mathcal{A} , \mathcal{B} and \mathcal{C}) have been developed [16] and have also provided the foundations for several successful approaches to automated planning. However, the main focus of these research efforts has been about reasoning within single agent domains. In single agent domains reasoning about information change mainly involves reasoning about what the agent knows about the world and how she can manipulate it to reach particular states. In multi-agent domains, on the other hand, an agent's action may change the world, other agents' knowledge about it and their knowledge about other agents' knowledge about the world. Similarly, goals of an agent in a multi-agent domain may involve manipulating the knowledge of other agents—in particular, this may concern not just their knowledge about the world, but also their knowledge about other agents' knowledge about the world.

As said before, epistemic planning is not only interested in the state of the world but also in the beliefs and the knowledge of the agents. Although there is a large body of research on multi-agent planning [7, 8, 18], very few efforts address the above mentioned aspects of multi-agent domains, which pose a number of new research challenges in representing and reasoning about actions and change.

¹In our case the relation between the world in a Kripke structure.

Due to its complexity, the majority of search based epistemic planners (e.g., [6, 10, 19, 20, 24, 29]) impose certain restrictions, such as the finiteness of the levels of nested beliefs. Such restrictions permit, for instance, in [19, 24, 29] to solve the problem by translating it into classical planning. Moreover, existing epistemic action languages [1, 2, 4, 12, 24] are able to model several families of problems, but cannot comprehensively reason on aspects like trust, dishonesty, deception and incomplete knowledge. In order to exploit epistemic reasoning in complex real-world scenarios it is paramount to define formalism that are able to consider the aforementioned concepts.

The planners [4, 12, 22] and the language $m\mathcal{A}^*$, that is the language that implemented by the these solvers, are the starting points of this research work. As we will explain in the following sections the main objectives of this thesis are: i) to formalize a complete and flexible epistemic action-based language; ii) to study alternative representations for epistemic models; and iii) to implement a competitive planner to reason on information change. Given the amount of information to properly introduce $m\mathcal{A}^*$ and the concept of *event update semantics* we redirect the reader who is not familiar with these two topics to [1] for a clear and complete explanation.

3 Goals of the research

The presented research aims to develop an epistemic action language and an epistemic solver general enough that could be adopted to explore the idea of comprehensive epistemic planning. In particular, this work aims to develop an environment in which it is possible to define custom *update models* and, therefore, *custom action types* in order to reason on all the subtle concepts related to the epistemology: e.g., trust, misconception, lies and so on.

Some existing epistemic solvers [12, 22] strictly relate action types to their underlying action language and, therefore, “statically” define the epistemic states update for each action type (i.e., *ontic*, *sensing* and *announcement*). These languages define precisely how an epistemic state (e-state) is updated after the execution of an action by defining how the e-states representation (i.e., *Kripke structure* or *possibilities*) is structurally updated. This update is done thanks to concepts of *agents’ observability* that allows to define the degree of change in the agents’ beliefs w.r.t. the properties of the world modified by the action. The newly generated e-state will then represents the “new” knowledge/beliefs graph that will allow checking belief formulae entailment without limitations (e.g., on the depth of the formulae). This approach permits to reason on the full extent of $\mathbf{S5}_n^C$ and $\mathbf{KD45}_n^C$ (with $n \geq 1$) but imposes limits on what an agent can do by allowing only predefined action types.

Other approaches, e.g. [24], are not based on structural e-state update and allow each action to define its effects explicitly. This means that each action needs to completely specify how the agents’ beliefs are updated. While having explicit updates allows to freely characterize each action’s effects, we feel that leaving such detail as input is partially in contrast with the idea of autonomous epistemic reasoning. Furthermore, it is easy to inadvertently omit or mis-characterize “intricate” knowledge/belief relations that should instead be handled by the autonomous solving process. Chains of knowledge/belief derived by different degree of observability and, as well as other complex knowledge/belief structures, may be easily mis-represented given their complex and not intuitive nature.

It is our objective to define an epistemic environment in which it is possible to define custom action types without having to rely on manually entering all the intricate knowledge/belief relations.

To be more precise about our research goals let us now present the major sub-goals of our work separately:

- To formalize a general and unified **action-based language** for epistemic domains [14]. The diver-

sification of approaches in solving this problem certainly brings benefits to the entire community of multi-agent epistemic planning. Nonetheless, with many specialized planners suitable for specific scenarios, the lack of a standardized way to define input problems has led to the creation of ad-hoc languages (each one with non-transferable feature, *e.g.*, explicit vs. implicit belief update) that can become an obstacle when comparing, combining (*i.e.*, ensemble) or testing different planners. We believe that a unified way of expressing MEP problems would provide researchers with a faster and less error-prone way of defining standardized benchmarks and would also relieve them from the burden of having to learn a new language each time they need to make use of a different tool.

- To define an **efficient e-state underlying data structure** and to formalize a **general e-state update** that is able to capture all the subtle possible variations of action types in epistemic reasoning. In particular, we envisioned two possible solutions: i) to enhance the actions types presented in by Fabiano *et al.* [12] so that they could capture the ideas of an *agent's attitude* and, therefore, reason on ideas such trust and misconception; and ii) to allow the user to define custom update models (and custom observability groups) so that each action could be associated to a specific way of structurally modifying the e-state. In particular, the latter solution seems the one that is more general and also more in line with the epistemic community.
- Finally, the most important contribution of this thesis is the implementation of an **general and comprehensive epistemic solver** that can be adopted by the community as a basis for future researches. This planner will be based on the PDDL-like epistemic action language so that researchers from other planning community can better understand how to define epistemic problems. Moreover, with the introduction of agents' attitudes and custom event model the planner will allow the users to tailor action in whichever fashion they prefer without having to worry about tedious and intricate effects definitions. This will help in formalizing new tools in which agents are able to reason while considering knowledge/belief relations with concepts such as lies, misconceptions, trust and so on. It would even be possible to define special groups of agents that reacts differently to the actions.

Whilst the generality of the planner is of the utmost importance, reducing the search times, given the inherent complexity of DEL, is also a feature that is essential to our solver. That is why, along with formalization of more efficient data structures and processes of e-state updates we also decided to focus on developing some domain-independent heuristics. These heuristics range from very simple ones, *e.g.*, the number of satisfied sub-goals, to more complex. An example of the latter is the updated version of the epistemic planning graph, that stems from a combination of [22] and [4].

Furthermore, this planner will be accompanied by a simple user interface in which the user will be able to select between several configurations; *e.g.*, whether to employ e-state reduction, to decrease generality in favor of performances, to print out the graphic representation of the planning process or even to select between different e-states representations. The interface will be further enriched with a *portfolio algorithm-like* procedure that will help the user to select the best planner configuration given a certain problem or even among the planner itself and some other solvers derived from the literature that may be more efficient in certain situations. In particular, the selection between the diverse planners, will be done following the ideas of [3] where meta-cognitive architectures are employed to select the best way of solving specific problems.

4 Status of the research

In this section we quickly present the work done on the goals introduced previously. Let us now provide, without tedious technical details, the status of the research of the goals presented in Section 3; the reader who is interested in a more complete introduction to these is mainly addressed to [3, 4, 12–14, 22].

- **Language formalization.** Considering that the main objective of this thesis is to provide a general and comprehensive epistemic solver, the first step of this process is to define an action language that will be the base for such solver. The formalization of this language followed diverse directions. Let us now briefly explain the various processes that allowed us to expand the language and to define a general and “standard” epistemic language.
 - The initial step of this research line was to adapt the language $m\mathcal{A}^*$ [1] so that it can be adopted also on a data structure that differs from the one considered by Baral *et al.*. In particular, while $m\mathcal{A}^*$ was defined on Kripke structures, we devised $m\mathcal{A}^P$ [12] that works on Possibilities [17].
 - Following we devised for both $m\mathcal{A}^*$ and $m\mathcal{A}^P$ the concept of *agents’ attitudes*. Thanks to this formalism we were able to add variations of the epistemic actions (in a paper under blind review) where the idea of trust and misconception is considered. This resulted in more complicated action types that allow to describe more realistic scenarios. Nevertheless, these actions types were still static and predetermined making the language incapable of handling all the nuances that naturally arise when concepts such as trust and lies are considered.
 - While trying to make the language more general we also encountered the problem of having a standardized way of defining domains across the epistemic community. We, therefore, decided to construct an input language for epistemic domains that could accommodate all the needs of the epistemic planning community while inheriting the syntax from the well known PDDL. From these considerations we created E-PDDL [14], an adaptation of PDDL for epistemic problems. We also created a parser [14] that is able to read a input in E-PDDL and convert such input into a format that the two state-of-the-art epistemic planners, *i.e.* [12, 24], can “understand”.
 - As a final step we envisioned a way of defining *custom update models* for both Kripke structure and Possibilities. These custom update models will allow the user to specify all sorts of behaviors for the actions making it possible to capture all the variations in the knowledge/belief update in epistemology. In particular, this level of customization will allow to confront several theories on how the agents’ knowledge/belief must be updated when, for example, in presence of lies, stubbornness, trust ignorance and so on. The specification of custom update models will also be made through a PDDL-like syntax.
- **Alternative e-state representation.** Kripke structure is the classical way to provide semantics for epistemic logic [15] but this does not mean that is the only suitable representation. In particular, in [12] we introduced a refactor of $m\mathcal{A}^*$, called $m\mathcal{A}^P$, that bases the state representation on *possibilities*, a data structure introduced in [17] and derived from *non-well-founded sets* theory. These structures allowed the planner to have an increase in terms of performances especially thanks to the optimized transition function. While this transition function is not general enough to accommodate the general case (*i.e.*, when action types are not statically defined) it still has good performances on the case when the needed action types are the “classical” ones; *i.e.*, ontic, sensing and announcement. Given the performance boost seen in [12], we decided to use possibilities as the main representation for e-states in the solver while maintaining also Kripke structures as it is the standard representation.
- **Epistemic Solver:** as main objective of our research we aim to provide a multi-agent epistemic planner flexible enough to be used in real world scenarios. This solver stems from the one firstly presented

in [22]. During the Ph.D. period the planner has been constantly modified in order to increase its performances and to include the new studied ideas. At this point of our study we have a modular C++ solver that:

- is able to “understand” input specification expressed in both $m\mathcal{A}^P$ and E-PDDL;
- can choose between two epistemic state representations, *i.e.*, Kripke structure or Possibilities;
- is able to adopt three different transition functions for each state representation: i) the one that only allows for the standard epistemic actions [1, 12]; ii) the update that considers the agents’ attitudes and is able to consider ideas such as trust and lies; and iii) the configuration where the action types can be defined by the user through custom update models. These configurations are all available as they represent different trade-offs between generality and performances;
- allows the use of heuristics. In particular, a refined version of the epistemic planning graph presented in [22] is under development. This planning graph will also be able to adapt to the selected type of transition function and it is independent of the chosen e-state representation making it available for every planner configuration.

Finally, a version in Answer Set Programming (ASP) of the planner in [12] has been implemented and presented in [4]. This planner has the same abilities as the one in [12] but thanks to its declarative nature it is more accessible to non-experts and easier to maintain. We are currently working in order to add the more general transition function also in the ASP solver. Both the versions (C++ and ASP) of the solver are available upon request.

5 Open issues

Multi-agent epistemic planning is a relatively new field of research in computer science. While the community is evolving rapidly there are still critical points that need to be addressed. In what follows we summarize what, we think, are the most critical issues that need to be tackled in epistemic planning.

- The lack of real-world examples where epistemic planning should be adopted. Multi-agent planners are very powerful tools able to reason on multi-agent scenarios, which are found everywhere in the real world. Nevertheless, the MEP community is missing strong examples that can be used as fundamental motivation of why such technology must be investigated. This lack of examples is probably due to the fact that epistemic reasoners do not scale well yet and, therefore, it is difficult to solve complex multi-agent problems when even toys examples are an obstacle. Another motivation for the absence of real-world inspired domains may be that each epistemic planner adopted its own specification language often making it difficult for other researches to adopt these tools in real tasks.
- As just mentioned another important issues is the poor scalability of the epistemic reasoners. Being the community relatively new it is normal that most of the effort is put in investigating the foundation of the problem rather than to optimize what already exists. Nonetheless, having tools that, most of the time, have not acceptable performances (w.r.t. classical planing, for example) limits the proliferation of the solvers themselves. It is paramount, in our opinion, to focus on the optimization of existing tools and formalism in order to have competitive epistemic reasoners that can be employed by other researchers or even in real-world scenarios. This would allow the community to gain more momentum and to grow even faster. Such optimizations could derive from several directions: implementation of heuristics, use of parallelism, adoption of symbolic e-state representations, combination of existing approach in a portfolio-like architecture and so on.

6 Acknowledgments

This research is partially supported by: the NSF HRD 1914635 and NSF HRD 1345232 projects, the University of Udine PRID ENCASE project and the GNCS-INdAM 2017–2020 projects.

References

- [1] Chitta Baral, Gregory Gelfond, Enrico Pontelli & Tran Cao Son (2015): *An Action Language for Multi-Agent Domains: Foundations*. CoRR abs/1511.01960. Available at <http://arxiv.org/abs/1511.01960>.
- [2] Thomas Bolander & Mikkel Birkegaard Andersen (2011): *Epistemic planning for single-and multi-agent systems*. *Journal of Applied Non-Classical Logics* 21(1), pp. 9–34, doi:10.1016/0010-0277(83)90004-5.
- [3] Grady Booch, Francesco Fabiano, Lior Horesh, Kiran Kate, Jonathan Lenchner, Nick Linck, Andrea Loreggia, Keerthiram Murugesan, Nicholas Mattei, Francesca Rossi & Biplav Srivastava (2021): *Thinking Fast and Slow in AI*. In: *Thirty-Fifth AAAI Conference on Artificial Intelligence, Virtual Event, AAAI Press*, pp. 15042–15046. Available at <https://ojs.aaai.org/index.php/AAAI/article/view/17765>.
- [4] Alessandro Burigana, Francesco Fabiano, Agostino Dovier & Enrico Pontelli (2020): *Modelling Multi-Agent Epistemic Planning in ASP*. *Theory and Practice of Logic Programming* 20(5), p. 593–608, doi:10.1017/S1471068420000289.
- [5] Alexander V. Chagrov & Michael Zakharyashev (1997): *Modal Logic*. *Oxford logic guides* 35, Oxford University Press.
- [6] Matthew Crosby, Anders Jonsson & Michael Rovatsos (2014): *A Single-Agent Approach to Multiagent Planning*. In: *ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic*, 263, IOS Press, pp. 237–242, doi:10.3233/978-1-61499-419-0-237.
- [7] Mathijs De Weerd, André Bos, Hans Tonino & Cees Witteveen (2003): *A resource logic for multi-agent plan merging*. *Annals of Mathematics and Artificial Intelligence* 37(1-2), pp. 93–130, doi:10.1023/A:1020236119243.
- [8] Mathijs De Weerd & Brad Clement (2009): *Introduction to planning in multiagent systems*. *Multiagent and Grid Systems* 5(4), pp. 345–355, doi:10.3233/MGS-2009-0133.
- [9] Agostino Dovier, Andrea Formisano & Enrico Pontelli (2013): *Autonomous agents coordination: Action languages meet CLP() and Linda*. *Theory and Practice of Logic Programming* 13(2), pp. 149–173, doi:10.1016/S0004-3702(00)00031-X.
- [10] Thorsten Engesser, Thomas Bolander, Robert Mattmüller & Bernhard Nebel (2017): *Cooperative Epistemic Multi-Agent Planning for Implicit Coordination*. In: *Proceedings of the Ninth Workshop on Methods for Modalities, M4M@ICLA 2017, Indian Institute of Technology, Kanpur, India, 8th to 10th January 2017, EPTCS* 243, pp. 75–90, doi:10.4204/EPTCS.243.6.
- [11] Francesco Fabiano (2019): *Design of a Solver for Multi-Agent Epistemic Planning*. In: *Proceedings 35th International Conference on Logic Programming (Technical Communications)*, pp. 403–412, doi:10.4204/EPTCS.306.54.
- [12] Francesco Fabiano, Alessandro Burigana, Agostino Dovier & Enrico Pontelli (2020): *EFP 2.0: A Multi-Agent Epistemic Solver with Multiple E-State Representations*. In: *Proceedings of the Thirtieth International Conference on Automated Planning and Scheduling, Nancy, France, October 26-30, 2020, AAAI Press*, pp. 101–109. Available at <https://aaai.org/ojs/index.php/ICAPS/article/view/6650>.
- [13] Francesco Fabiano, Idriss Riouak, Agostino Dovier & Enrico Pontelli (2019): *Non-Well-Founded Set Based Multi-Agent Epistemic Action Language*. In: *Proceedings of the 34th Italian Conference on Computational Logic, Trieste, Italy, June 19-21, 2019, CEUR Workshop Proceedings* 2396, CEUR-WS.org, pp. 242–259. Available at <http://ceur-ws.org/Vol-2396/paper38.pdf>.

- [14] Francesco Fabiano, Biplav Srivastava, Jonathan Lenchner, Lior Horesh, Francesca Rossi & Marianna Bergamaschi Ganapini (2021): *E-PDDL: A Standardized Way of Defining Epistemic Planning Problems*. In: *Knowledge Engineering for Planning and Scheduling*, Online, p. in press. Available at https://icaps21.icaps-conference.org/workshops/KEPS/Papers/KEPS_2021_paper_3.pdf.
- [15] Ronald Fagin & Joseph Y Halpern (1994): *Reasoning about knowledge and probability*. *Journal of the ACM (JACM)* 41(2), pp. 340–367, doi:10.1145/174652.174658.
- [16] Michael Gelfond & Vladimir Lifschitz (1998): *Action Languages*. *Electron. Trans. Artif. Intell.* 2, pp. 193–210. Available at <http://www.ep.liu.se/ej/etai/1998/007/>.
- [17] J. Gerbrandy & W. Groeneveld (1997): *Reasoning about information change*. *Journal of Logic, Language and Information* 6(2), pp. 147–169, doi:10.1023/A:1008222603071.
- [18] Claudia V Goldman & Shlomo Zilberstein (2004): *Decentralized control of cooperative systems: Categorization and complexity analysis*. *J. Artif. Intell. Res.(JAIR)* 22, pp. 143–174, doi:10.1613/jair.1427.
- [19] Xiao Huang, Biqing Fang, Hai Wan & Yongmei Liu (2017): *A General Multi-agent Epistemic Planner Based on Higher-order Belief Change*. In: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, ijcai.org, pp. 1093–1101, doi:10.24963/ijcai.2017/152.
- [20] Filippos Kominis & Hector Geffner (2017): *Multiagent Online Planning with Nested Beliefs and Dialogue*. In: *Proceedings of the International Conference on Automated Planning and Scheduling, ICAPS, Pittsburgh, Pennsylvania, USA*, pp. 186–194. Available at <https://aaai.org/ocs/index.php/ICAPS/ICAPS17/paper/view/15748>.
- [21] Saul A. Kripke (1963): *Semantical Analysis of Modal Logic I Normal Modal Propositional Calculi*. *Mathematical Logic Quarterly* 9(5-6), pp. 67–96, doi:10.1002/malq.19630090502.
- [22] Tiep Le, Francesco Fabiano, Tran Cao Son & Enrico Pontelli (2018): *EFP and PG-EFP: Epistemic Forward Search Planners in Multi-Agent Domains*. In: *Proceedings of the Twenty-Eighth International Conference on Automated Planning and Scheduling, AAAI Press, Delft, The Netherlands*, pp. 161–170. Available at <https://aaai.org/ocs/index.php/ICAPS/ICAPS18/paper/view/17733>.
- [23] Nir Lipovetzky & Hector Geffner (2017): *Best-First Width Search: Exploration and Exploitation in Classical Planning*. In: *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, San Francisco, California, USA*, pp. 3590–3596. Available at <http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14862>.
- [24] Christian J. Muise, Vaishak Belle, Paolo Felli, Sheila A. McIlraith, Tim Miller, Adrian R. Pearce & Liz Sonenberg (2015): *Planning Over Multi-Agent Epistemic States: A Classical Planning Approach*. In: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*, AAAI Press, pp. 3327–3334. Available at <http://www.aaai.org/ocs/index.php/AAAI/AAAI15/paper/view/9974>.
- [25] Silvia Richter & Matthias Westphal (2010): *The LAMA planner: Guiding cost-based anytime planning with landmarks*. *Journal of Artificial Intelligence Research* 39, pp. 127–177, doi:10.1613/jair.2972.
- [26] Stuart J. Russell & Peter Norvig (2010): *Artificial Intelligence - A Modern Approach, Third International Edition*. Pearson Education. Available at http://vig.pearsoned.com/store/product/1,1207,store-12521_isbn-0136042597,00.html.
- [27] Raymond R. Smullyan (1968): *First-order logic*. 43, Springer-Verlag Berlin Heidelberg, doi:10.1007/978-3-642-86718-7.
- [28] Hans Van Ditmarsch, Wiebe van Der Hoek & Barteld Kooi (2007): *Dynamic epistemic logic*. 337, Springer Netherlands, doi:10.1007/978-1-4020-5839-4.
- [29] Hai Wan, Rui Yang, Liangda Fang, Yongmei Liu & Huada Xu (2015): *A Complete Epistemic Planner without the Epistemic Closed World Assumption*. In: *IJCAI International Joint Conference on Artificial Intelligence, Buenos Aires, Argentina*, pp. 3257–3263. Available at <http://ijcai.org/Abstract/15/459>.