A Simple Framework for Cognitive Planning (Extended Abstract)

Jorge Fernandez Davila, Dominique Longin, Emiliano Lorini and Frédéric Maris

IRIT-CNRS, Toulouse University, France

Abstract

We present a model of cognitive planning that was published in the proceedings of AAAI-2021 [1]. The model generalizes epistemic planning. It has been recently implemented in a conversational agent and applied to the domain of AI-based coaching.

1. Introduction

Classical planning in artificial intelligence (AI) is the general problem of finding a sequence of actions (or operations) aimed at achieving a certain goal [2]. It has been shown that classical planning can be expressed in the propositional logic setting whereby the goal to be achieved is represented by a propositional formula [3]. In recent times, epistemic planning was proposed as a generalization of classical planning in which the goal to be achieved can be epistemic [4, 5]. For example, in epistemic planning, the planning agent could try to reveal a secret to the target agent 1 thereby making agent 1 know the secret, while keeping the target agent 2 uninformed. This requires the use of more expressive languages that allow to represent epistemic attitudes such as knowledge and belief. The standard languages for epistemic planning are epistemic logic (EL) [6] and dynamic epistemic logic (DEL), the dynamic extension of EL by so-called event models [7]. A variety of epistemic logic languages with different levels of expressivity and complexity have been introduced to formally represent the epistemic planning problem and to efficiently automate it (see, *e.g.*, [8, 9, 10, 11, 12, 13]).

2. From Epistemic to Cognitive Planning

In a recent paper [1], we introduced cognitive planning as a further generalization of epistemic planning. We formalized it in the epistemic logic framework presented in [14]. Unlike the standard semantics for EL using multi-relational Kripke models, the semantics presented in [14] use belief bases. It has been recently applied to modeling multi-agent belief revision [15].

In cognitive planning, it is not only some knowledge or belief state of a target agent that is to be achieved, but more generally a cognitive state. The latter could involve not only knowledge and beliefs, but also goals, intentions and, more generally, motivations. Cognitive planning makes clear the distinction between persuasion (i.e., inducing someone to believe that a certain

CAKR'22: Cognitive Aspects of Knowledge Representation, IJCAI-ECAI 2022 Workshop, July 23–29, 2022, Vienna, Austria

CEUR Workshop Proceedings (CEUR-WS.org)

fact is true) and influence (i.e., motivating someone to behave in a certain way) and elucidates the connection between these two notions. Specifically, since beliefs are the input of decision-making and provide reasons for deciding and for acting, the persuader can indirectly change the persuadee's motivations and behaviors by changing her beliefs, through the execution of a sequence of speech acts. In other words, in cognitive planning, the persuader could try to modify the persuadee's beliefs in order to affect persuadee's motivations. Moreover, cognitive planning takes resource boundedness and limited rationality of the persuadee seriously. For this reason, it is particularly well-suited for human-machine interaction (HMI) applications in which an artificial agent is expected to interact with a human — who is by definition resource-bounded — through dialogue and to induce her to behave in a certain way. These two aspects are exemplified in Figure 1. The artificial agent has both (i) a model of the human's overall cognitive state, and (ii) a persuading or influencing goal towards the human. Given (i) and (ii), it tries to find a sequence of speech acts aimed at modifying the human's cognitive state thereby guaranteeing the achievement of its persuading/influencing goal.

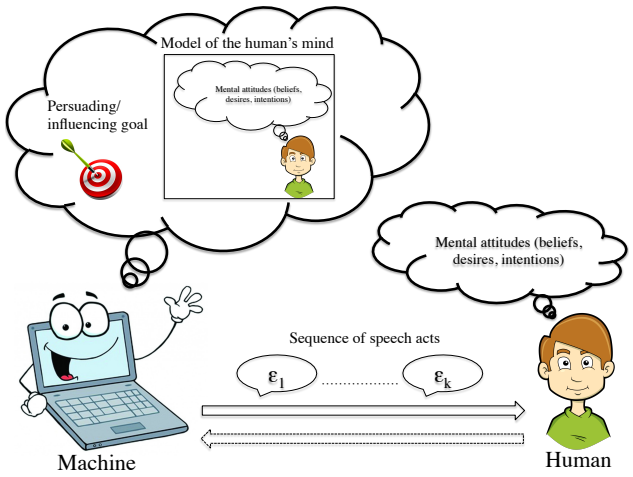

Figure 1: Cognitive planning

Models of persuasion in AI are mostly based on argumentation. (See [16] for a general introduction to the research in this area.) Some of these models are built on Walton & Krabbe's notion of persuasion dialogue in which one party seeks to persuade another party to adopt a belief or point-of-view she does not currently hold [17]. There exist models based on abstract argumentation [18, 19, 20, 21] as well probabilistic models where the persuader's uncertainty about what the persuadee knows or believes is represented [22]. There exist also models based on possibility theory in which a piece of information is represented as an argument which can be more or less accepted depending on the trustworthiness of the agent who proposes it [23]. More recently, argumentation-based models of planning for persuasion have been proposed in which actions in a plan are abstract arguments [21]. In our approach beliefs of the persuader and of the persuadee are explicitly modeled. Moreover, the components of a plan are speech acts either of type *assertion* or of type *question* with a specific logical content.

3. Application to Conversational Agents

The model and algorithm of cognitive planning presented in [1] were recently implemented in a artificial agent (see [24] for more details). In [25] we successfully applied our cognitive planning approach to modelling conversational agents in the human-machine interaction (HMI) domain. In particular, we developed an artificial coaching system based on motivational interviewing, a counseling method used in clinical psychology for eliciting behavior change [26].

Acknowledgments

This work is supported by the ANR project CoPains ("Cognitive Planning in Persuasive Multimodal Communication").

References

[1] J. Fernandez, D. Longin, E. Lorini, F. Maris, A simple framework for cognitive planning, in: Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-2021), AAAI Press, 2021, pp. 6331–6339. URL: https://ojs.aaai.org/index.php/AAAI/article/view/16786.

[2] M. Ghallab, D. Nau, P. Traverso, Automated planning: theory and practice, Morgan Kaufmann, 2004.

[3] T. Bylander, The computational complexity of propositional STRIPS planning, Artificial Intelligence 69 (1994) 165–204.

[4] T. Bolander, M. B. Andersen, Epistemic planning for single- and multi-agent systems, Journal of Applied Non-Classical Logics 21 (2011) 9–34.

[5] B. Löwe, E. Pacuit, A. Witzel, DEL planning and some tractable cases, in: Proceedings of the 3rd International International Workshop on Logic, Rationality and Interaction (LORI 2011), Springer Berlin Heidelberg, 2011, pp. 179–192.

[6] J. Y. Halpern, Y. Moses, A guide to completeness and complexity for modal logics of knowledge and belief, Artificial Intelligence 54 (1992) 319–379.

[7] H. P. van Ditmarsch, W. van der Hoek, B. Kooi, Dynamic Epistemic Logic, Kluwer Academic Publishers, 2007.

[8] C. Muise, V. Belle, P. Felli, S. A. McIlraith, T. Miller, A. R. Pearce, L. Sonenberg, Planning over multi-agent epistemic states: A classical planning approach, in: AAAI 2015, AAAI Press, 2015, pp. 3327–3334.

[9] C. Muise, V. Belle, P. Felli, S. A. McIlraith, T. Miller, A. R. Pearce, , L. Sonenberg, Efficient multi-agent epistemic planning: Teaching planners about nested belief, Artificial Intelligence 302 (2021).

[10] F. Kominis, H. Geffner, Beliefs in multiagent planning: from one agent to many, in: ICAPS 2015, AAAI Press, 2015, pp. 147–155.

[11] M. C. Cooper, A. Herzig, F. Maffre, F. Maris, E. Perrotin, P. Régnier, A lightweight epistemic logic and its application to planning, Artificial Intelligence 298 (2021). URL: https://doi.org/10.1016/j.artint.2020.103437. doi:10.1016/j.artint.2020.103437.

[12] C. Baral, G. Gelfond, E. Pontelli, T. C. Son, An action language for multi-agent domains, Artificial Intelligence 302 (2022) 103601. URL: https://www.sciencedirect.com/science/article/pii/S0004370221001521. doi:https://doi.org/10.1016/j.artint.2021.103601.

[13] A. Burigana, F. Fabiano, A. Dovier, E. Pontelli, Modelling multi-agent epistemic planning in asp, Theory and Practice of Logic Programming 20 (2020) 593–608. doi:10.1017/S1471068420000289.

[14] E. Lorini, Rethinking epistemic logic with belief bases, Artificial Intelligence 282 (2020). doi:https://doi.org/10.1016/j.artint.2020.103233.

[15] E. Lorini, F. Schwarzentruber, Multi-agent belief base revision, in: Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI 2021), ijcai.org, 2021.

[16] H. Prakken, Formal systems for persuasion dialogue, The Knowledge Engineering Review 21 (2006) 163–188.

[17] D. Walton, E. Krabbe, Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning, SUNY Series in Logic and Language, State University of New York Press, 1995.

[18] T. J. M. Bench-Capon, Persuasion in practical argument using value-based argumentation frameworks, Journal of Logic and Computation 13(3) (2003) 429–448.

[19] E. Bonzon, N. Maudet, On the outcomes of multiparty persuasion, in: Proceedings of the 8th International Conference on Argumentation in Multi-Agent Systems (ArgMAS 2011), Springer-Verlag, 2011, p. 86–101.

[20] L. Amgoud, N. Maudet, S. Parsons, Modelling dialogues using argumentation, in: Proceedings of the Fourth International Conference on MultiAgent Systems, IEEE, 2000, pp. 31–38.

[21] E. Black, A. J. Coles, C. Hampson, Planning for persuasion, in: Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), volume 2, IFAAMAS, 2017, pp. 933–942.

[22] A. Hunter, Modelling the persuadee in asymmetric argumentation dialogues for persuasion, in: Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI 2015), AAAI Press, 2015, p. 3055–3061.

[23] C. Da Costa Pereira, A. Tettamanzi, S. Villata, Changing one's mind: Erase or rewind? possibilistic belief revision with fuzzy argumentation based on trust, in: Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI 2011), AAAI Press, 2011, p. 164–171.

[24] J. Fernandez, D. Longin, E. Lorini, F. Maris, An implemented system for cognitive planning, in: A. P. Rocha, L. Steels, H. J. van den Herik (Eds.), Proceedings of the 14th International Conference on Agents and Artificial Intelligence, ICAART 2022, Volume 3, Online Streaming, February 3-5, 2022, SCITEPRESS, 2022, pp. 492–499.

[25] E. Lorini, N. Sabouret, B. Ravenet, J. Fernandez, C. Clavel, Cognitive planning in motivational interviewing, in: A. P. Rocha, L. Steels, H. J. van den Herik (Eds.), Proceedings of the 14th International Conference on Agents and Artificial Intelligence, ICAART 2022, Volume 2, Online Streaming, February 3-5, 2022, SCITEPRESS, 2022, pp. 508–517.

[26] B. Lundahl, B. L. Burke, The effectiveness and applicability of motivational interviewing: A practice-friendly review of four meta-analyses, Journal of clinical psychology 65 (2009) 1232–1245.