An Algebraic Structuring of Epistemic States for BDI Agents in Uncertain Environments

Doctoral Consortium

Charles A. N. Costa
University of Brasília
Brasília, Brazil
charles.costa@aluno.unb.br

ABSTRACT

Autonomous agents operating in dynamic and partially observable environments must reason and plan under uncertainty. Standard Belief–Desire–Intention (BDI) architectures typically rely on binary belief representations, which limit their ability to model uncertain information. Extensions such as CAN+ address this limitation by introducing graded beliefs represented as weighted formulas. However, these representations lack strong algebraic properties, restricting systematic belief manipulation. This paper proposes an algebraic structuring of epistemic states for BDI agents, modeling them as an Abelian group. This structure preserves the expressiveness of graded beliefs while enabling algebraic reasoning techniques, such as cancellation and equation solving. The approach is illustrated through plan executability analysis in CAN+, demonstrating how algebraic properties improve reasoning about uncertain beliefs during plan selection and execution.

KEYWORDS

Belief-Desire-Intention, epistemic states, uncertain beliefs, algebra

ACM Reference Format:
Charles A. N. Costa. 2026. An Algebraic Structuring of Epistemic States for BDI Agents in Uncertain Environments: Doctoral Consortium. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/HAWU4840

1 INTRODUCTION

Autonomous agents operating in dynamic and partially observable environments must plan and act under uncertainty. In the domain of Belief-Desire-Intention (BDI) agents, the information gathered from the environment constitutes the beliefs of the agent [4]. Traditional languages for the development of BDI agents, such as AgentSpeak [3] and the Conceptual Notation Agent (CAN) [10], do not explicitly support the modelling of uncertain beliefs. However, there are proposals to provide languages and environments to fill this gap. One of those approaches is CAN+ [1, 2, 11], a language that enhances the CAN language by adding support for epistemic beliefs.

CAN+ represents the agent's beliefs by sets of weighted formulas, where the formula is a possible state of the world, and the weight is a numerical value that corresponds to the degree of confidence the agent associates with the belief. Bauters et al. [1] introduced a revision operator that enables the agent to revise its beliefs in response to changes in the state of the world.

Unfortunately, a disadvantage of CAN+ is the difficulty of manipulating epistemic states and incorporating them into existing algorithms. To overcome this caveat, we propose an algebraic treatment of epistemic states, introducing modifications to the CAN+ set of postulates so that epistemic states and their revision operator could be treated as Abelian groups [6]. In Abelian groups, operations are reversible, and the order of application does not change the output. One of these modifications introduces the concept of the inverse of epistemic beliefs, which enables algorithms based on backward search on solution spaces.

Problem Statement. The central problem addressed in this work is the lack of algebraic structure in graded belief representations used by BDI agents. In particular, the absence of inverses makes it difficult to analyze the combined effects of belief updates.

The research question guiding this work is: *How can epistemic states representing uncertain beliefs be structured to improve both expressiveness and reasoning capabilities in BDI agents?* The hypothesis is that structuring epistemic states as an Abelian group provides a principled foundation for belief manipulation while remaining compatible with existing BDI semantics.

Contributions. The main contributions of this research are advancements in the formal representation of epistemic states within BDI agents. This research demonstrates how this algebraically structured approach not only enhances belief manipulability but also preserves the expressive power required for complex decision-making processes.

2 ALGEBRAIC STRUCTURING OF EPISTEMIC STATES

A typical formalization of epistemic states is Ordinal Conditional Functions (OFCs) by Spohn [12], which associate an integer value indicating a degree of plausibility with a possible world. The OFCs are the basis of Ma and Liu [7] formalization of epistemic states for managing uncertain beliefs, which, in turn, are the basis for Bauters et al. [2] definition of the language CAN+. According to Masolo and Porello [8], a weighted formula is a logical formula, typically in propositional logic, associated with a weight, usually a real number, that expresses the degree to which the formula contributes to the representation of a concept.

In CAN+, the extensional representation of a weighted belief base is given by $\phi = \{(\omega, w) \mid \omega \in \mathcal{L}_\phi \text{ and } w \in \mathbb{Z} \cup \{\infty, -\infty\}\}$. Where $\mathcal{L}_\phi$ is the language of ϕ, ω is a formula in $\mathcal{L}_\phi$, and w is

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/HAWU4840

the degree of confidence of ω. It is also required that no pair of formulas $(\omega_1, \omega_2) \in \phi$ have any model in common. The epistemic valuation λ induced by ϕ is given by $\lambda(\alpha) = max\{\Phi(\omega) \mid \omega \models \alpha\}$. We can now define our belief revision operator as follows:

Definition 2.1. Let G be the set of all weighted belief bases, $\mathcal{A}, \mathcal{B} \in G$. We define the operator $\oplus_s : G \times G \mapsto G$ as $\mathcal{A} \oplus_s \mathcal{B} = \mathcal{A}^+_{\mathcal{B}} \cup \mathcal{A}^-_{\mathcal{B}} \cup \mathcal{B}^-_{\mathcal{A}}$, where:

$$\mathcal{A}^+_{\mathcal{B}} = \{(\alpha \wedge \beta, m + n) \mid (\alpha, m) \in \mathcal{A} \text{ and } (\beta, n) \in \mathcal{B}\}$$
$$\mathcal{A}^-_{\mathcal{B}} = \{(\alpha \wedge \neg \bigvee(\mathcal{B}^*), m) \mid (\alpha, m) \in \mathcal{A}\}$$
$$\mathcal{B}^-_{\mathcal{A}} = \{(\beta \wedge \neg \bigvee(\mathcal{A}^*), n) \mid (\beta, n) \in \mathcal{B}\}$$

Having $C^* = \{\varphi \mid (\varphi, k) \in C\}$, the set of formulas without their weights. The sets above omit all pairs (φ, k) s.t. φ is inconsistent.

Definition 2.1 generalizes the syntactic revision operator of [2], considering the intuition behind it, i.e., the result of the revision of a weighted belief base by another consists of updating the credence of all shared models between them, while maintaining the other models untouched. Figure 1 presents this idea, where Ω is the set of all possible worlds, $\mathcal{A}$, $\mathcal{B}$ are weighted belief bases, and the gray circles represent the models of the formulas in $\mathcal{A}$ and $\mathcal{B}$.

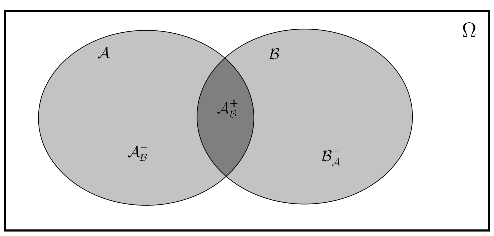

Figure 1: Venn diagram of the models of two weighted belief bases, A for $\mathcal{A}$ and B for $\mathcal{B}$. The dark gray area represents the result of the revision of $\mathcal{A}$ by $\mathcal{B}$.

Using a carefully crafted set of postulates, we were able to prove that $[\mathcal{A}] \oplus_s [\mathcal{B}] = [\mathcal{A} \oplus_s \mathcal{B}]$ is an abelian group. In our proposal, it is worth noting that the inverse of a weighted belief base $\mathcal{A}$ is given by $-\mathcal{A} = \{(\omega, -w) \mid (\omega, w) \in \mathcal{A}\}$, and $\mathcal{A} \oplus_s -\mathcal{A} = \emptyset$. Having said that, $\emptyset$ is the identity element. This structuring allows us to use cancellation of terms to solve equations involving weighted belief sets, such as the following example:

Example 2.2. The weighted belief base of an agent is $\mathcal{B}_0 = (\alpha, 10)$. We have $\lambda(a) = 10$ and $\lambda(\neg a) = 0$. Thus, since $\lambda(\alpha) > \lambda(\neg\alpha)$, $\mathcal{B} \models \alpha$. The agent desires to bring about a possible world $\mathcal{B}_1$ such that $\mathcal{B}_1 \models \neg\alpha$. What would be the input that results in this state of the world? First, let us say that the goal is to make $\mathcal{B}_1 = \{(\alpha, 10), (\neg\alpha, 20)\}$. The agent can use our proposal to reason as follows:

$$\mathcal{B} \oplus_s x = \mathcal{B}_1$$
$$\{(\alpha, 10)\} \oplus_s x = \{(\alpha, 10), (\neg\alpha, 20)\}$$
$$\{(\alpha, 10)\} \oplus_s x \oplus_s \{(\alpha, -10)\} = \{(\alpha, 10), (\neg\alpha, 20)\} \oplus_s \{(\alpha, -10)\}$$
$$x = \{(\neg\alpha, 20)\}$$

3 APPLICATION TO AUTOMATED PLANNING

In traditional automated planning, an action scheme a is the triple $\langle Pre(a), Add(a), Del(a)\rangle$. $Pre(a)$ are the preconditions that must be met to execute the action. $Add(a)$ and $Del(a)$ are the effects that set the literals true or false, respectively. This representation supports regressive reasoning, i.e., computing the state that must precede the execution of an action. In CAN+ [2, 11], the action's preconditions are formulas in the epistemic state languages. The effects are sequences of inputs since actions can only change the state of the world by revision. CAN+, thus, defines an action scheme a as $\langle Pre(a), Pos(a)\rangle$. In our approach, this representation supports regressive reasoning by the application of inverses.

Let us take the case of the backward search algorithm for automated planning. The backward search is a classical automated planning algorithm. It starts from the goal G and iteratively regresses to the initial state [9]. At each step, the algorithm selects a relevant action from the agent's repertoire of actions. An action is relevant if $G \cap Add(a) \neq \emptyset$ and $G \cap Del(a) = \emptyset$. It then computes a new regressed goal G' by $G' = (G \setminus Add(a)) \cup Pre(a)$. Applying our algebraic approach, we calculate the regression goal G' using $G' = G \oplus_s (-Pos(a)) \oplus_s Pre(a)$. Only taking that the relevant actions are those in which, for all literals $G \models l$, we have $Pos(a) \models l$.

We can apply the same reasoning to the calculus of plan executability. As said in Yao et al. [13], plans have hidden conditions which derive from the sequence of actions and parallel compositions. Determining these conditions at compilation time helps with scheduling intentions that will not block each other's execution. In traditional hierarchical planning, this is not a trivial task. In CAN+, calculating these conditions becomes even harder. Applying our framework, deriving the plan conditions can be softened by the application of combinations and simplifications allowed by group operations.

4 OPEN QUESTIONS

An open question is the computational cost of operating with weighted belief states. Bauters et al. [2] state that determination whether $\lambda(\phi) = m$ is NP-complete since it reduces from the SAT problem. However, tractable versions of the problem can mitigate this problem. Bauters et al. [2] propose a tractable epistemic state restricting new information to weighted literals. We intend to investigate whether we could admit wider conditions, such as a limited number of literals or new information restricted to the Disjunctive Normal Form (DNF). Another open question is the conflict between the revision operator and the AGM postulates [5], since our revision operator is symmetric. We aim to investigate whether the conflict with AGM postulates arises from our revision operator or from the uncertainty of the environment. Finally, an illustrative application that more closely reflects a real-world problem has yet to be developed.

ACKNOWLEDGMENTS

I would like to thank my advisor, Célia Ghedini Ralha, for her guidance and advice. This material is based on work accepted to be published in the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026).

REFERENCES

[1] Kim Bauters, Weiru Liu, Jun Hong, Carles Sierra, and Lluís Godo. 2014. CAN(PLAN)+: extending the operational semantics of the BDI architecture to deal with uncertain information. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence* (Quebec City, Quebec, Canada) *(UAI'14)*. AUAI Press, Arlington, Virginia, USA, 52–61.

[2] Kim Bauters, Kevin McAreavey, Weiru Liu, Jun Hong, Lluis Godo, and Carles Sierra. 2017. Managing different sources of uncertainty in a BDI framework in a principled way with tractable fragments. *Journal of Artificial Intelligence Research* 58 (2017), 731–775.

[3] Rafael H Bordini, Jomi Fred Hübner, and Michael Wooldridge. 2007. *Programming multi-agent systems in AgentSpeak using Jason*. Vol. 15. John Wiley & Sons, Chichester, UK.

[4] Michael Bratman. 1987. *Intention, Plans, and Practical Reason.* Cambridge, MA: Harvard University Press, Cambridge.

[5] Giorgos Flouris, Dimitris Plexousakis, and Grigoris Antoniou. 2004. Generalizing the AGM postulates: preliminary results and applications.. In *NMR*. 171–179.

[6] Thomas W Hungerford. 2012. *Algebra*. Vol. 73. Springer Science & Business Media, Heidelberg, Germany.

[7] Jianbing Ma and Weiru Liu. 2011. A framework for managing uncertain inputs: An axiomization of rewarding. *International Journal of Approximate Reasoning* 52, 7 (2011), 917–934.

[8] Claudio Masolo and Daniele Porello. 2018. Representing concepts by weighted formulas. In *Formal Ontology in Information Systems*. IOS Press, 55–68.

[9] Stuart J. Russell and Peter Norvig. 2009. *Artificial Intelligence: a modern approach* (3 ed.). Pearson, London, England.

[10] Sebastian Sardina, Lavindra De Silva, and Lin Padgham. 2006. Hierarchical planning in BDI agent programming languages: A formal approach. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*. IFAAMAS, Japan, 1001–1008.

[11] Aidan Scannell. 2018. *Extending BDI agents to model and reason with uncertainty*. Master of Research in Robotics and Autonomous Systems. University of Bristol and University of the West of England.

[12] Wolfgang Spohn. 1988. Ordinal conditional functions: A dynamic theory of epistemic states. In *Causation in decision, belief change, and statistics: Proceedings of the Irvine Conference on Probability and Causation*. Springer, Heidelberg, Germany, 105–134.

[13] Yuan Yao, Lavindra de Silva, and Brian Logan. 2016. Reasoning about the executability of goal-plan trees. In *Engineering Multi-Agent Systems: 4th International Workshop, EMAS 2016, Singapore, Singapore, May 9-10, 2016, Revised, Selected, and Invited Papers 4*. Springer, 176–191.