

Empowering BDI Agents with Generalised Decision-Making

Blue Sky Ideas Track

Ramon Fraga Pereira
University of Manchester
Manchester, England, United Kingdom
ramon.fragapereira@manchester.ac.uk

Felipe Meneguzzi
University of Aberdeen
Aberdeen, Scotland, United Kingdom
felipe.meneguzzi@abdn.ac.uk

ABSTRACT

While research on software agents has long focused on explicit agent communication, there is comparatively less effort on implicit communication between agents via recognising each other’s intentions and desires for understanding their decision-making reasoning process. Since most human communication is not explicit, we aim to outline a research agenda to help endow autonomous agents with analogous coordination capabilities. In this paper, we formalise a framework that empowers the decision-making process of BDI agents in adversarial and cooperative environments by casting them as generalised planners using *Theory of Mind*. Our formalisation uses the fundamental philosophical properties of the BDI model and its reasoning process to outline a broad research agenda in agents’ research.

ACM Reference Format:

Ramon Fraga Pereira and Felipe Meneguzzi. 2023. Empowering BDI Agents with Generalised Decision-Making: Blue Sky Ideas Track. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 5 pages.

1 INTRODUCTION

Research on autonomous agents has long been concerned with reasoning about an agent’s actions within an environment [35], regardless of the underlying agent architecture [15]. Indeed, the agent architecture based on *Beliefs*, *Desires*, and *Intentions* (BDI) [5], explicitly distinguishes the kind of reasoning about the state of the world (theoretical reasoning) from reasoning towards actions (practical reasoning). This separation of reasoning concerns allows BDI-type agents to engage in a type of introspective reasoning about an agent’s goals and its relation with feasibility as part of the agent’s reasoning cycle. *Automated Planning* [20], by contrast, reasons about how sequential decisions about actions lead to desired outcomes over time. Being able to formally model an agent’s internal mechanism not only allows an individual agent to reason about its own goals, it allows agents to build models of other agents in the environment. Such capability creates the possibility of improving an agent’s interaction with others by anticipating each other’s behaviours, minimising the need for explicit communication. Thus, since planning is a key capability for autonomous agent architectures [15, 35, 42], we argue for renewed impetus on the study of the interplay between agent and planning formalisms.

While recent work on planning has started to bridge the gap between automated planning and agents [42], such work often ignores the contributions from the agent’s research altogether. One key area between research on agent systems and on planning lies in generalised planning [25] as a means-ends reasoning mechanism. *Generalised Planning* is a type of planning problem that encodes multiple different individual goals, the solution of which offers maximal coverage. Finding such generalised plans corresponds to intention-selection in BDI agents with declarative goals [35].

This paper lays the foundations of formal underpinnings for BDI agents as generalised planners, paving the way for a number of new research directions on the interaction between planning techniques and agent reasoning in two key ways. First, our formalisation of generalised planning as BDI means-ends reasoning allows an agent to reason about its own goals (desires) as a generalised planning problem. Such reasoning about goals is a critical requirement for agents driven by declarative goals, since it allows agents to reason about contingencies for agent synthesis [3]. Second, our formalisation of *Intent (Goal) Recognition* in the generalised planning setting using *Theory of Mind* [46], together with our first contribution, allows agents to reason and understand the intentions of other agents [24, 61]. This also allows us to cast the problem of understanding the decision-making of autonomous agents, employing the concept of *Theory of Mind* to allow our generalised BDI agents to simulate an approximate recognition process in order to estimate the behaviour of other agents, as it would be able to identify if the other agents in the environment are acting in an adversarial or cooperative way. Our vision for research on BDI agents provides an ambitious road-map for research in the agents community, leading to a number of challenges, which we refine throughout the paper.

2 BACKGROUND AND RELATED WORK

2.1 Planning

The *environment* in which our autonomous agents act and operate for achieving their goals follows the formalism of *Classical Planning*. Here, a *domain model* representing the environment is *fully observable*, has *discrete* properties, and the actions’ outcomes (i.e., effects) are *deterministic* [20, Chapter 2].

A *planning domain model* Ξ is a $\langle \mathcal{F}, \mathcal{A} \rangle$ where: \mathcal{F} is a set of *fluents* (i.e., environment properties); and \mathcal{A} is a set of *actions* where every action $a \in \mathcal{A}$ has a *positive cost*, denoted as $cost(a)$, and its own set of preconditions, add and delete lists: $Pre(a), Add(a), Del(a)$. A state S is a finite set of positive fluents $f \in \mathcal{F}$ that follows the *closed world assumption* so that if $f \in S$, then f is true in S . We assume a simple inference relation \models such that $S \models f$ iff $f \in S$, $S \not\models f$ iff $f \notin S$, and $S \models f_0 \wedge \dots \wedge f_n$ iff $\{f_0, \dots, f_n\} \subseteq S$. An action $a \in \mathcal{A}$ is

applicable to a state S if and only if $S \models \text{Pre}(a)$, and generates a new state S' such that $S' \leftarrow (S \cup \text{Add}(a))/\text{Del}(a)$.

A *planning problem* \mathcal{P} is a $\langle \Xi, s_0, s_g \rangle$ where: Ξ is a planning domain as described above; $s_0 \subseteq \mathcal{F}$ is the *initial state*; and $s_g \subseteq \mathcal{F}$ is a *goal state*. A *solution* to the planning problem \mathcal{P} is a *plan* $\pi = [a_0, \dots, a_n]$ that maps s_0 into a state $S \models s_g$, that is, in which the goal state s_g holds. The cost of a plan $\pi = [a_0, a_1, \dots, a_n]$ is $\text{cost}(\pi) = \sum \text{cost}(a_i)$; we say that a plan π^* is *optimal* if there exists no other plan π' for \mathcal{P} such that $\text{cost}(\pi') < \text{cost}(\pi^*)$. The main purpose of a planning algorithm $\text{PLANNER}(\Xi, s_0, s_g)$ is to find such plans. We say a planner is *optimising* if it guarantees to find one of the optimal plans, and *satisficing* otherwise.

2.2 Generalised Planning

We define *generalised planning problems* following [8, 25]. A *generalised planning problem* \mathcal{GP} is defined as $\langle \mathcal{P}_0, \mathcal{P}_1, \dots, \mathcal{P}_N \rangle$, and represents the problem of solving a set of planning problems that share some common structure (also known as sharing the *agent*), i.e., the action scheme and the environment properties. Namely, a *generalised planning problem* is a finite set of N *planning problems*, so we define a *generalised planning problem* \mathcal{GP} with $N \geq 2$. A *solution* to a *generalised planning problem* \mathcal{GP} is a *generalised plan* Π that solves a \mathcal{GP} . We define $\text{exec}(\Pi, \mathcal{P}) = [a_0, \dots, a_n]$ as an analogy to a plan π , which is a sequence of actions that solves a planning problem \mathcal{P} . Thus, for every $\mathcal{P}_i \in \mathcal{GP}, 1 \leq i \leq N$, $\text{exec}(\Pi, \mathcal{P}_i)$ solves \mathcal{P}_i . Like Section 2.1, we call $\mathcal{GPLANNER}(\mathcal{GP})$ an algorithm that solves such problems. Recent work develops different types of *generalised planning* approaches (i.e., *planners*), such as *generalised planning* approaches that rely on off-the-shelf *Classical Planning* [54], *heuristic search* [55, 57], and *landmarks* [56].

2.3 Goal Recognition as Planning

Goal Recognition is the task of discerning the intended goal agent aims to achieve given a sequence of observations, whereas *Plan Recognition* consists of identifying/recognising the plan (i.e., sequence of actions) that achieves such intended goal [41]. Model-based *Goal Recognition* (also known as *Goal and Plan Recognition as Planning*) was formally defined by Ramirez and Geffner in [49, 50], where they formally define the task of recognising goals and plans over a *planning domain theory*. Such formalism allows specifying agents' possible behaviours using action schemes and declarative goals, and it usually follows well-known planning formalisms in the literature, such as STRIPS [7] and PDDL [34].

A *recognition problem* \mathcal{P}_{GR} [37] is a tuple $\langle \Xi, s_0, G, \Omega_\pi \rangle$ where: $\Xi = \langle \mathcal{F}, \mathcal{A} \rangle$ is *planning domain*, \mathcal{F} is a set of fluents, and \mathcal{A} is a set of actions; s_0 is the initial state; G is the set of possible goals, which includes the *correct intended goal*¹ s_g^* (i.e., $s_g^* \in G$), and $\Omega_\pi = [o_1, \dots, o_n]$ is a sequence of observations $o_i \in \mathcal{A}$. Ω_π represents a valid plan (i.e., achievable, given the initial state s_0), and it usually represents a partial sequence of actions (with possibly missing observed actions) that aims to achieve s_g^* .

The *solution* to a *recognition problem* \mathcal{P}_{GR} is a goal (or a set of goals) which the recognition approach determines to be the correct intended goal (s). As a way to interpret the recognition process and rank/categorise the goal(s) that the observed agents (most likely)

¹Note that the *correct intended goal* s_g^* is **unknown** for the recogniser.

aim to achieve [33], existing recognition approaches often return either a *probability distribution* over the set of goals [50, 63], or *scores* associated with each goal in the set of possible goals [17, 43, 52]. In this paper, recognition techniques return a *probability distribution* over the possible agent intentions. Existing model-based recognition approaches rely on a range of techniques, such as *planning* and *adapted heuristic functions* [49, 50], *planning graphs* [17], *top-k* and *diverse planning* [63], *landmarks* [43, 62], *learning* and *symbolic planning* [1, 44, 67], *linear programming* [52] and *goal/plan mirroring* [64, 65], and *multiple-goal heuristic search* [19].

3 BDI AGENTS AS GENERALISED PLANNERS

Research on BDI agents has yielded a number of formalisms of varying complexities [5, 6, 15]. These often cater to agent programming languages such as AgentSpeak(L) [51]. We focus on the more general type of BDI architecture, in which the agent designer only specifies *declarative goals* [14] coupled with a particular environment specification that follows a STRIPS-style [18] planning formalism². We can formalise this architecture at a very basic level as a tuple $\langle \mathcal{B}, \mathcal{D}, \mathcal{I} \rangle$, such that each element of the tuple represents the key components of a BDI agent, as illustrated in Figure 1. Beliefs




Figure 1: BDI Agent as Generalised Planning.

\mathcal{B} represent the agent's internal representation of the state of the environment. Much like the states in planning domains, beliefs correspond to a set of fluents \mathcal{F} , which the agent updates via a sensing function. This corresponds to the initial state from which the agent plans and compares the results of its actions. Desires \mathcal{D} correspond to the set of **potential** goals G that the agent can (but does not necessarily) pursue at any given time. In our formalisation, desires are sets of tuples $\langle \varphi_i, D_i, \sigma_i \rangle$, containing conjunctive formulas representing potential goals (D_i), as well as a context condition (φ_i) indicating when a goal becomes relevant, and a preference value σ . Such relevance condition helps the agent filter intentions and facilitates connecting BDI-style agents to generalised planners. Indeed, the set of desires is analogous to a set of generalised planning problems \mathcal{GP} , which we leverage next. The preference ordering σ allows an agent to prioritise certain goals. To simplify notation, when all goals have the same priority, we represent a desire tuple as $\langle \varphi_i, D_i \rangle$, omitting the priority. Since desires need not be internally consistent, the agent's reasoning cycle selects a subset of desires toward the achievement of which it commits to. Once the agent commits to such desires, these become the intentions \mathcal{I} , towards which the agent plans. An intention consists of a tuple $\langle \langle \varphi_j, D_j \rangle, \pi_j \rangle$ comprising a desire and a plan to which the agent has committed. Intentions are desires and associated plans towards

²Our vision is not limited to deterministic and fully observable environments, since the reasoning process relies on off-the-shelf solvers.

Algorithm 1 Abstract BDI Reasoning Cycle

```

1: procedure REASONINGCYCLE( $\mathcal{B}, \mathcal{D}, \mathcal{I}, \Xi$ )
2:   loop
3:      $\mathcal{B} \leftarrow \mathcal{B} \cup \text{SENSE}()$ 
4:     while  $\mathcal{I}$  is not empty do
5:       Pick an intention  $\langle \langle \varphi, D \rangle, \pi \rangle \in \mathcal{I}$  s.t.  $\mathcal{B} \models \varphi \wedge \neg D$ 
6:       ACT( $\pi$ )
7:       Find  $\{ \langle \varphi_1, D_1 \rangle \dots \langle \varphi_n, D_n \rangle \} \in \mathcal{D}^2$ 
         s.t.  $\exists \Pi, \Pi = \mathcal{G}\text{PLANNER}(\{ \langle \Xi, \mathcal{B}, D_1 \rangle \dots \langle \Xi, \mathcal{B}, D_n \rangle \})$ 
8:        $\mathcal{I} \leftarrow \{ \langle \langle \varphi_1, D_1 \rangle, \Pi \rangle, \langle \langle \varphi_n, D_n \rangle, \Pi \rangle \}$ 

```

which the agent commits to act. At each point in time, the agent then can run a planning process using $\langle \Xi, \mathcal{B}, \mathcal{I} \rangle$ as its problem.

From these components, we can define a simplified reasoning cycle in the pseudocode of Algorithm 1 that encapsulates the basic conceptual operation of a BDI agent. In this case, the agent uses a *generalised planning algorithm* as the means-ends reasoner for an agent driven entirely by declarative goals. This reasoning cycle corresponds to the conceptual work of Bratman [5] more closely than common implementations of BDI programming languages [13, 27, 51]. While we are aware of no contemporary implementation of BDI agents using this kind of reasoning cycle, we posit that current progress on planning technology and AI in general will yield novel, more powerful BDI interpreters along these lines. Ultimately, this reasoning cycle allows one to reason about an observed agent’s intentions using the formal machinery of Section 2.3.

This formalisation has a number of desirable properties within our vision. First, generalised planning naturally represents that other agents in the environment might have multiple goals in parallel. This is entirely independent of a plan library, thus relaxing the assumption that the local recogniser knows about the plan libraries of other agents in competitive settings. Second, it is agnostic to either adversarial or cooperative settings, as we elaborate further.

4 GENERALISED INTENT RECOGNITION

Using the basic machinery of a generic goal-driven BDI agent, we can formalise the problem of recognising such an agent’s intent by reasoning over its assumed desires and observed actions in an environment. Thus, we define a *generalised recognition problem* as $\langle \mathbb{G}, \Omega_\Pi \rangle$, where $\mathbb{G} = \langle \mathcal{G}\mathcal{P}_0, \mathcal{G}\mathcal{P}_1, \dots, \mathcal{G}\mathcal{P}_N \rangle$ is a set of *generalised planning problems* with different goal states (i.e., desires), and Ω_Π is a sequence of observations that represent a *generalised plan* Π to solve the intended generalised planning problem $\mathcal{G}\mathcal{P}^* \in \mathbb{G}$. Figure 2 illustrates the *generalised recognition process* when observing BDI agents as *generalised planners*. Following the recognition approaches in the literature, a recogniser takes as input a sequence of observations Ω_Π , representing a *generalised plan* Π that a BDI agent is executing to achieve its intentions \mathcal{I} filtered from the agent’s desires \mathcal{D} , returning a probability distribution over the set of possible agent’s desires (committed to as intentions), defined as \mathbb{G} .

Solving a *generalised recognition problem* consists of computing posterior probabilities over \mathbb{G} given Ω_Π . Inspired by the probabilistic formulation of Ramírez and Geffner for *Goal and Plan Recognition as Planning* in [50], we define a probability distribution for every possible intention $\mathcal{G}\mathcal{P}$ in the set of intentions \mathbb{G} , and the sequence

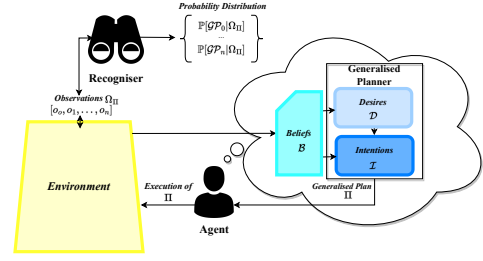


Figure 2: BDI Generalised Recognition process.

of observations Ω_Π to be the Bayesian posterior conditional probability, as follows:

$$\mathbb{P}(\mathcal{G}\mathcal{P} \mid \Omega_\Pi) = \eta * \mathbb{P}(\Omega_\Pi \mid \mathcal{G}\mathcal{P}) * \mathbb{P}(\mathcal{G}\mathcal{P}) \quad (1)$$

where $\mathbb{P}(\mathcal{G}\mathcal{P})$ is a *prior probability* to an intention $\mathcal{G}\mathcal{P}$, η is a *normalizing factor* ($\eta = [\sum_{\mathcal{G}\mathcal{P} \in \mathbb{G}} \mathbb{P}(\Omega_\Pi \mid \mathcal{G}\mathcal{P}) * \mathbb{P}(\mathcal{G}\mathcal{P})]^{-1}$), and $\mathbb{P}(\Omega_\Pi \mid \mathcal{G}\mathcal{P})$ is the probability of observing Ω_Π when the intention is $\mathcal{G}\mathcal{P}$. $\mathbb{P}(\Omega_\Pi \mid \mathcal{G}\mathcal{P})$ expresses the probability of observing Ω_Π when the BDI agent’s intention is $\mathcal{G}\mathcal{P}$. This formulation allows us to incorporate knowledge of an agent’s preferences over desires as the prior probability $\mathbb{P}(\mathcal{G}\mathcal{P})$: the more preferred a desired state is, the more likely the agent is to select it as an intention. When the *priors* are uniform (e.g., no information about preferences), the most likely intention(s) are precisely the ones that maximise $\mathbb{P}(\Omega_\Pi \mid \mathcal{G}\mathcal{P})$. Namely, with equal priors, the higher $\mathbb{P}(\Omega_\Pi \mid \mathcal{G}\mathcal{P})$ is, the likelier the desire $\mathcal{G}\mathcal{P}$ is to be part of the intentions of the observed BDI agent. Solving this problem efficiently creates a substantial challenge for research in agent systems.

5 APPROACHES AND CHALLENGES

With our visionary ideas of *Generalised BDI Agents* and *Generalised Intent Recognition* in place, we lay out different avenues for further research and promising approaches in the decision-making for BDI agents. This includes a discussion of hurdles and potential solutions to deal with adversarial and cooperative contexts/settings, as well as links to research on BDI agents, *AI Planning*, and *Intent Recognition*.

Within our formalisation, the key initial challenge is computing $\mathbb{P}(\Omega_\Pi \mid \mathcal{G}\mathcal{P})$, which is key to recognising other agents’ intentions.

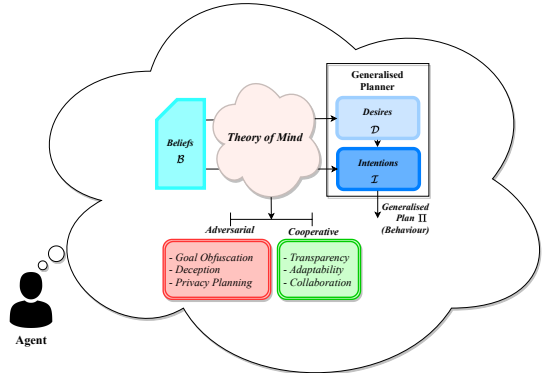


Figure 3: BDI agent employing *Theory of Mind*.

In line with current recognition approaches, one can compute what we call as *general recognition score* ϵ . We define the *general recognition score* ϵ as way to compute $\mathbb{P}(\Omega_{\Pi} \mid \mathcal{GP})$ in Equation 2.

$$\mathbb{P}(\Omega_{\Pi} \mid \mathcal{GP}) = [1 + (1 - \epsilon)]^{-1} \quad (2)$$

The *general recognition score* ϵ is a value between 0 and 1, i.e., $0 \leq \epsilon \leq 1$, and represents how compliant the sequence of observations Ω_{Π} is to a generalised plan Π for achieving \mathcal{GP} . The closer ϵ is to 1, the greater $\mathbb{P}(\Omega_{\Pi} \mid \mathcal{GP})$ is, and therefore, the more likely \mathcal{GP} is to be the intention that the observed agent BDI aims to achieve. Whereas the closer ϵ is to 0, the lower $\mathbb{P}(\Omega_{\Pi} \mid \mathcal{GP})$ is, and thus, less likely \mathcal{GP} is to be the current intention.

Given the effectiveness of seminal and recent approaches to *Goal and Plan Recognition as Planning*, the computation of the *general recognition score* ϵ could be done by adapting well-known recognition approaches in the literature [43, 50, 64], making this type of recognition process more general.

5.1 Promising Approaches

The ability to perform goal recognition within our *Generalised BDI* agents constitutes a practical *Theory of Mind* as part of their decision-making process. This is an effective mechanism for an agent to reason over other agents' actions to predict the actions, desires, and intentions of others acting in the same environment.

There are two key ways in which generalised BDI agents can make their decisions: either taking into account, or ignoring other agents' behaviour. First, Decision-Making that is **agnostic (or independent)** of other agents is often the norm in BDI agent research. This makes their decision-making process based solely on their own beliefs, desires, and intentions without considering the larger context of the environment or the presence of other agents [51]. In this case, our generalised BDI agents would follow the reasoning cycle presented in Section 3, also shown in Figure 1, having their decision-making process general, in contrast to existing BDI agent reasoning processes [26, 38]. Second, decision-making that is **fully aware** of other agents is critical for BDI agents that make decisions considering cooperative [10, 21], adversarial [48], or both types of agents (or environment contexts/settings) [29]. Such awareness of other agents, in turn, entails understanding two different stances towards other agents in the environment.

In a **cooperative context**, agents endowed with *Generalised Intent Recognition* as part of their *Theory of Mind* [24, 59–61] can identify when other agents are acting in a cooperative way. The interactions (e.g., via executions of plans) of other agents acting in an environment allow our recognition approach to return a probability distribution over their possible goals (desires), and if their intended goals (desires) align with the goals of our BDI agent, then it may indicate that the other agents are acting cooperatively. Therefore, in order to make the decision-making of our BDI agent cooperative, more transparent and clearer, we envision a decision-making reasoning process in which the agents deliberately plan to achieve their goals using a *Cooperative Planning* technique as part of its *Generalised Planning* process, such *Transparent* [31], *Legible* [9, 16], or *Adaptable* planning processes (Figure 3, *Theory of Mind* – light-yellow box, *Cooperative*).

By contrast, in an **adversarial context**, a BDI agent may have to act privately for self-preservation [30, 58], to make their actions

as private as possible, obfuscating other agents [2]. Alternatively, agents can deliberately deceive other agents due to strategic purposes for protecting their goals [53], surprise and ambush other agents, or thwart their intentions [45]. Thus, to facilitate adversarial decision-making within our BDI agent, we envision using *Theory of Mind* approaches [24, 59–61] to recognise and comprehend the desires of adversarial agents, and then turn the generalised planning process of our BDI plans consciously in a non-cooperative way, using non-cooperative planning techniques as part of their decision-making process, such as deception [47], privacy planning [28], etc, see (Figure 3, *Theory of Mind* – light-red box, *Adversarial*).

5.2 Challenges and Opportunities

While the potential approaches to reasoning about other agents outlined in Section 5.1 are ambitious in pushing agents research ahead, they also entail various challenges. First, planning problems in general and generalised planning in particular is computationally hard. Thus, in practice, the entire reasoning cycle with planning and recognition might be costly in terms of computation time. Second, one of the key assumptions of virtually all goal recognition approaches is knowledge of the goal hypothesis space. Thus, the challenge remains of recognising other agents' goals when the possible goals are unknown. Agents' goals might change over time and they will certainly not inform the observers when changing their goals, especially in adversarial (non-cooperative) environment settings. Third, how to determine the level of rationality of the other agents' behaviour to have successful and efficient reasoning and decision-making cycles. We argue that the *rationality measure* of Masters and Sardiña [32] could be applied as part of the recognition process to deal with both rational and irrational agents. Fourth, our simplified reasoning cycle abstracts away issues of dealing deal with failures, desire reconsideration, and re-planning. This necessitates further progress in the design of planning-driven, purely declarative agent interpreters. Fifth, predicting the actions of an agent with multiple objectives is hard (as shown in [12]).

Here, our vision of generalised planning paves the way to do that, by attempting to recognise general desires (sets of goals). Finally, while recognising the behaviour of other agents by observation has been studied in the context of inferring social norms [11, 36, 40], and its violations [39], our richer formalisation of the problem and its solutions should also expand research on normative reasoning.

6 DISCUSSION

This paper offers a research vision for *Intent Recognition* over BDI agents capable of *Generalised Planning*. To the best of our knowledge, only Xu et al. [66] attempts to formalising intent recognition in the context of BDI agents. This approach, however, uses a very limited notion of BDI agency, that of BDI agents driven by plan libraries. Our research vision, by contrast, provides a more complex conceptualisation of BDI agents, which is arguably more faithful to the philosophical formalisation of BDI [5], and in line with recent advances in BDI architectures [4]. Such a conceptualisation is technically possible as a result of substantial recent progress on automated planning [22, 23, 25]. Nevertheless, fully incorporating such progress into agent architectures requires substantial work on BDI agent architectures that can effectively use these capabilities.

REFERENCES

- [1] Leonardo Amado, Ramon Fraga Pereira, Joao Paulo Aires, Mauricio Magnaguagno, Roger Granada, and Felipe Meneguzzi. 2018. Goal Recognition in Latent Space. In *IJCNN*.
- [2] Sara Bernardini, Fabio Fagnani, and Santiago Franco. 2020. An Optimization Approach to Robust Goal Obfuscation. In *KR*.
- [3] Blai Bonet, Giuseppe De Giacomo, Hector Geffner, Fabio Patrizi, and Sasha Rubin. 2020. High-level Programming via Generalized Planning and LTL Synthesis. In *KR*.
- [4] Rafael H. Bordini, Amal El Fallah Seghrouchni, Koen V. Hindriks, Brian Logan, and Alessandro Ricci. 2020. Agent programming in the cognitive era. *JAAMAS* 34, 2 (2020).
- [5] Michael E. Bratman. 1984. Two Faces of Intention. *Philosophical Review* 93 (1984), 375–405.
- [6] Michael E. Bratman, David J. Israel, and Martha E. Pollack. 1988. Plans and Resource-Bounded Practical Reasoning. *Computational Intelligence* (1988), 349–355.
- [7] Tom Bylander. 1994. The Computational Complexity of Propositional STRIPS Planning. *Journal of Artificial Intelligence Research (JAIR)* 69 (1994), 165–204.
- [8] Sergio Jiménez Celorrio, Javier Segovia-Aguas, and Anders Jonsson. 2019. A Review of Generalized Planning. *Knowledge Engineering Review* 34 (2019).
- [9] Elizabeth Cha, Anca D. Dragan, and Siddhartha S. Srinivasa. 2013. Effects of robot capability on user acceptance. In *HRI*.
- [10] Angela Consoli, Jeffrey Tweedale, and Lakhmi C. Jain. 2007. An Architecture for Agent Coordination and Cooperation. In *KES*.
- [11] Natalia Criado. 2018. Resource-bounded Norm Monitoring In Multi-agent Systems. *Journal of Artificial Intelligence Research* 62 (2018), 153–192.
- [12] Michael Dann, Yuan Yao, Natasha Alechina, Brian Logan, and John Thangarajah. 2022. Multi-Agent Intention Progression with Reward Machines. In *IJCAI*.
- [13] Mehdi Dastani. 2008. 2APL: A Practical Agent Programming Language. *JAAMAS* 16 (2008), 214–248.
- [14] Mehdi Dastani, M. Birna van Riemsdijk, and Michael Winikoff. 2011. Rich goal types in agent programming. In *AAMAS*.
- [15] Lavindra de Silva, Felipe Meneguzzi, and Brian Logan. 2020. BDI Agent Architectures: A Survey. In *IJCAI*.
- [16] Anca D. Dragan, Kenton C. T. Lee, and Siddhartha S. Srinivasa. 2013. Legibility and predictability of robot motion. In *HRI*.
- [17] Yolanda E.-Martín, María D. R.-Moreno, and David E. Smith. 2015. A fast goal recognition technique based on interaction estimates. In *IJCAI*.
- [18] Richard Fikes and Nils Nilsson. 1971. STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. *Artificial Intelligence* 2, 3-4 (1971).
- [19] Richard G. Freedman, Yi Ren Fung, Roman Ganchin, and Shlomo Zilberstein. 2018. Towards Quicker Probabilistic Recognition with Multiple Goal Heuristic Search. In *AAAI Workshop on Plan, Activity, and Intent Recognition (PAIR)*.
- [20] Hector Geffner and Blai Bonet. 2013. *A Concise Introduction to Models and Methods for Automated Planning*. Morgan & Claypool Publishers.
- [21] Nathan Griffiths and Michael Luck. 1999. Cooperative Plan Selection Through Trust. In *Workshop on Modelling Autonomous Agents in a Multi-Agent World*.
- [22] Malte Helmert. 2006. The Fast Downward Planning System. *Journal of Artificial Intelligence Research* 26 (2006), 191–246. <https://www.fast-downward.org>
- [23] Malte Helmert and Carmel Domshlak. 2009. Landmarks, Critical Paths and Abstractions: What’s the Difference Anyway?. In *ICAPS*.
- [24] Mark K. Ho, Rebecca Saxe, and Fiery Cushman. 2022. Planning with Theory of Mind. *Trends in Cognitive Sciences* 26, 11 (2022), 959–971.
- [25] Yuxiao Hu and Giuseppe De Giacomo. 2011. Generalized Planning: Synthesizing Plans that Work for Multiple Environments. In *IJCAI*.
- [26] Jomi F. Hubner, Jaime S. Sichman, and Olivier Boissier. 2007. Developing organised multiagent systems using the MOISE+ model: programming issues at the system and agent levels. *International Journal of Agent-Oriented Software Engineering* 1, 3/4 (2007), 370–395.
- [27] François Félix Ingrand and Vianney Coutance. 2001. Real-Time Reasoning Using Procedural Reasoning. Technical Report.
- [28] Sarah Keren, Avigdor Gal, and Erez Karpas. 2016. Privacy Preserving Plans in Partially Observable Environments. In *IJCAI*.
- [29] Anagha Kulkarni, Siddharth Srivastava, and Subbarao Kambhampati. 2019. A unified framework for planning in adversarial and cooperative environments. In *AAAI*.
- [30] Zhengshang Liu, Yue Yang, Tim Miller, and Peta Masters. 2021. Deceptive Reinforcement Learning for Privacy-Preserving Planning. In *AAMAS*.
- [31] Aleck M. MacNally, Nir Lipovetzky, Miquel Ramírez, and Adrian R. Pearce. 2018. Action Selection for Transparent Planning. In *AAMAS*.
- [32] Peta Masters and Sebastian Sardiña. 2021. Expecting the unexpected: Goal recognition for rational and irrational agents. *Artificial Intelligence* 297 (2021).
- [33] Peta Masters and Mor Vered. 2021. What’s the Context? Implicit and Explicit Assumptions in Model-Based Goal Recognition. In *IJCAI*.
- [34] Drew McDermott, Malik Ghallab, Adele Howe, Craig Knoblock, Ashwin Ram, Manuela Veloso, Daniel Weld, and David Wilkins. 1998. PDDL – The Planning Domain Definition Language. In *AI Planning and Scheduling Conference*.
- [35] Felipe Meneguzzi and Lavindra de Silva. 2015. Planning in BDI agents: a survey of the integration of planning algorithms and agent reasoning. *KER* 30, 1 (2015), 1–44.
- [36] Felipe Meneguzzi, Brian Logan, and Moser Silva Fagundes. 2014. Norm monitoring with asymmetric information. In *AAMAS*.
- [37] Felipe Meneguzzi and Ramon Fraga Pereira. 2021. A Survey on Goal Recognition as Planning. In *IJCAI*.
- [38] Felipe Meneguzzi, Odinaldo Rodrigues, Nir Oren, Wamberto Weber Vasconcelos, and Michael Luck. 2015. BDI Reasoning with Normative Considerations. *Engineering Applications Artificial Intelligence* 43 (2015), 127–146.
- [39] Felipe Meneguzzi, Odinaldo Rodrigues, Nir Oren, Wamberto W. Vasconcelos, and Michael Luck. 2015. BDI reasoning with normative considerations. *Engineering Applications of Artificial Intelligence* 43 (2015), 127–146.
- [40] Felipe Meneguzzi, Pankaj R. Telang, and Munindar P. Singh. 2013. A First-Order Formalization of Commitments and Goals for Planning. In *AAAI*.
- [41] Reuth Mirsky, Sarah Keren, and Christopher W. Geib. 2021. *Introduction to Symbolic Plan and Goal Recognition*. Morgan & Claypool Publishers.
- [42] Sunandita Patra, James Mason, Amit Kumar, Malik Ghallab, Paolo Traverso, and Dana Nau. 2021. Integrating Acting, Planning, and Learning in Hierarchical Operational Models. *Artificial Intelligence* 30 (2021), 478–487.
- [43] Ramon Fraga Pereira, Nir Oren, and Felipe Meneguzzi. 2020. Landmark-Based Approaches for Goal Recognition as Planning. *Artificial Intelligence* 279 (2020).
- [44] Ramon Fraga Pereira, Mor Vered, Felipe Meneguzzi, and Miquel Ramírez. 2019. Online Probabilistic Goal Recognition over Nominal Models. In *IJCAI*.
- [45] Alberto Pozanco, Yolanda E.-Martín, Susana Fernández, and Daniel Borrajo. 2018. Counterplanning using Goal Recognition and Landmarks. In *IJCAI*.
- [46] David Premack and Guy Woodruff. 1978. Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences* 1, 4 (1978), 515–526.
- [47] Adrian Price, Ramon Fraga Pereira, Peta Masters, and Mor Vered. 2023. Domain-Independent Deceptive Planning. In *AAMAS*.
- [48] Simon Pulawski, Hoa Khanh Dam, and Aditya Ghose. 2021. BDI-Dojo: developing robust BDI agents in evolving adversarial environments. In *ACSOS*.
- [49] Miquel Ramírez and Hector Geffner. 2009. Plan Recognition as Planning. In *IJCAI*.
- [50] Miguel Ramírez and Hector Geffner. 2010. Probabilistic Plan Recognition using off-the-shelf Classical Planners. In *AAAI*.
- [51] Anand S. Rao. 1996. AgentSpeak(L): BDI agents speak out in a logical computable language. In *Proceedings of the 7th European workshop on Modelling autonomous agents in a multi-agent world*, Vol. 1038. Springer-Verlag, 42–55.
- [52] Luisa R. de A. Santos, Felipe Meneguzzi, Ramon Fraga Pereira, and Andre Pereira. 2021. An LP-Based Approach for Goal Recognition as Planning. In *AAAI*.
- [53] Stefan Sarkadi, Alison R. Panisson, Rafael H. Bordini, Peter McBurney, Simon Parsons, and Martin Chapman. 2019. Modelling Deception using Theory of Mind in Multi-agent Systems. *AI Communications* 32, 4 (2019), 287–302.
- [54] Javier Segovia-Aguas, Sergio Jiménez Celorrio, and Anders Jonsson. 2019. Computing Programs for Generalized Planning using a Classical Planner. *Artificial Intelligence* 272 (2019), 52–85.
- [55] Javier Segovia-Aguas, Sergio Jiménez Celorrio, and Anders Jonsson. 2022. Computing Programs for Generalized Planning as Heuristic Search. In *IJCAI*.
- [56] Javier Segovia-Aguas, Sergio Jiménez Celorrio, Laura Sebastián, and Anders Jonsson. 2022. Scaling-Up Generalized Planning as Heuristic Search with Landmarks. In *SOCS*.
- [57] Javier Segovia-Aguas, Sergio Jiménez, and Anders Jonsson. 2021. Generalized Planning as Heuristic Search. In *ICAPS*.
- [58] Guy Shani. 2018. Advances and Challenges in Privacy Preserving Planning. In *IJCAI*.
- [59] Maayan Shvo, Ruthrash Hari, Ziggy O’Reilly, Sophia Abolore, Sze-Yuh Nina Wang, and Sheila A. McIlraith. 2022. Proactive Robotic Assistance via Theory of Mind. In *IROS*.
- [60] Maayan Shvo, Toryn Q. Klassen, and Sheila A. McIlraith. 2020. Towards the Role of Theory of Mind in Explanation. In *AAMAS Workshop on Explainable, Transparent Autonomous Agents and Multi-Agent Systems*.
- [61] Maayan Shvo, Toryn Q. Klassen, and Sheila A. McIlraith. 2022. Resolving Misconceptions about the Plans of Agents via Theory of Mind. In *ICAPS*.
- [62] Maayan Shvo and Sheila A. McIlraith. 2020. Active Goal Recognition. In *AAAI*.
- [63] Shirin Sohrabi, Anton V. Riabov, and Octavian Udrea. 2016. Plan Recognition as Planning Revisited. In *IJCAI*.
- [64] Mor Vered, Gal A. Kaminka, and Sivan Biham. 2016. Online Goal Recognition through Mirroring: Humans and Agents. In *ACS*.
- [65] Mor Vered, Ramon Fraga Pereira, Gal Kaminka, and Felipe Meneguzzi. 2018. Towards Online Goal Recognition combining Goal Mirroring and Landmarks. In *AAMAS*.
- [66] Hongyun Xu, Bastin Tony Roy Savarimuthu, Aditya Ghose, Evan D. Morrison, Qiying Cao, and Youqun Shi. 2013. Automatic BDI Plan Recognition from Process Execution Logs and Effect Logs. In *EMAS*.
- [67] Tan Zhi-Xuan, Jordyn L. Mann, Tom Silver, Josh Tenenbaum, and Vikash Mansinghka. 2020. Online Bayesian Goal Inference for Boundedly Rational Planning Agents. In *NeuRIPS*.