Empowering BDI Agents with Generalised Decision-Making

Ramon Fraga Pereira† **Felipe Meneguzzi‡**

†The University of Manchester, England, UK
ramon.fragapereira@manchester.ac.uk
‡University of Aberdeen, Scotland, UK
felipe.meneguzzi@abdn.ac.uk

Auckland, May 2024

Table of Contents

BDI Agents and Planning

What?

- Agent architecture based on three "mental" structures:
  - Beliefs, Desires, and Intentions
- Based on a philosophical model for *practical reasoning*
- Provides a blueprint for agent reasoning, suitable for:
  - Agent implementations
  - Reasoning about other agents
- Key process is means ends reasoning:
  - Typically using a *plan library*
  - More recent work focuses on *automated planning*

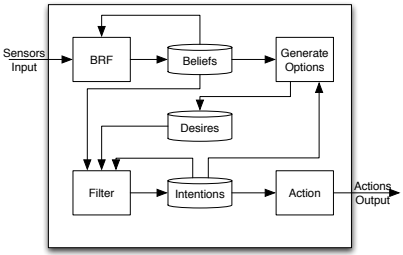

Planning in BDI

Why?

- Focus of much research in AAMAS for the past three decades, primarily, on:
  - Agent Oriented Software Engineering
  - Agent reasoning cycle
  - Multiagent systems (populated by BDI agents)
- Relatively fewer efforts on the interface of means-ends reasoning and the agent model

Table of Contents

Definition (Planning Task)

A planning task $\Pi = \langle \Xi, s_0, G \rangle$ is a tuple composed of a domain definition Ξ, an initial state s_0, and a goal state specification G. A solution to a planning task is a plan or policy π that reaches a goal state G starting from the initial state s_0 by following the transitions defined in the domain definition Ξ.

Planning problems have three key ingredients

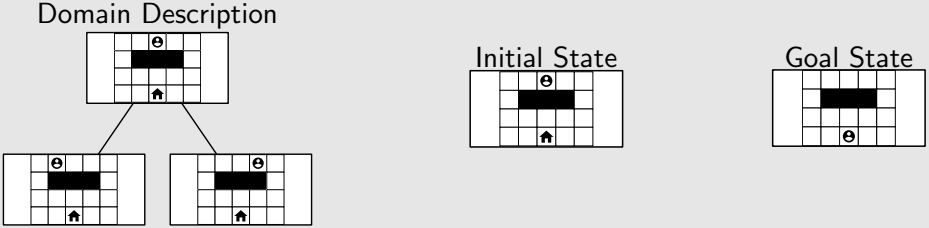

Solution

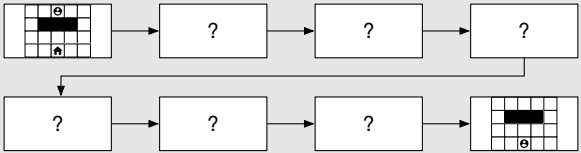

Planning problems have three key ingredients

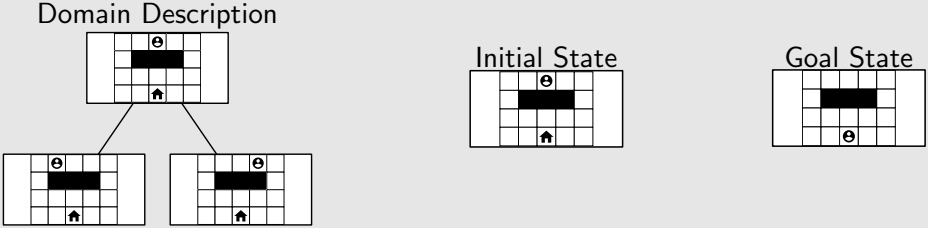

Solution

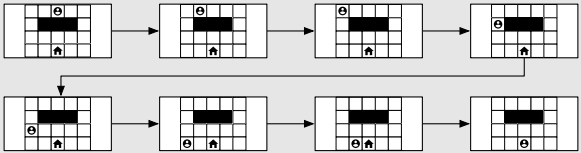

Definition (Goal Recognition Task)

A goal recognition task $\Pi_{\mathcal{G}}^{\Omega_\pi} = \langle \Xi, s_0, \mathcal{G}, \Omega_\pi \rangle$ is a tuple composed of a domain definition Ξ, an initial state s_0, a set of goal hypotheses $\mathcal{G}$, and a sequence of observations Ω_π.

Goal/Plan Recognition problems have **four** key ingredients

Domain Description

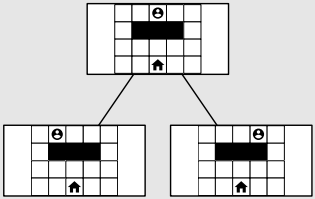

Initial State

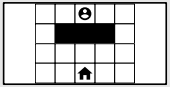

Goal **Hypotheses**

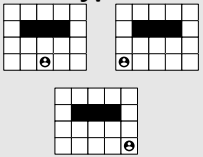

Observations

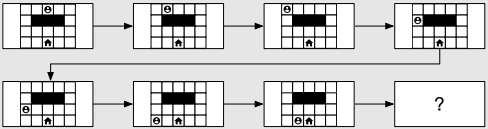

Goal/Plan Recognition problems have **four** key ingredients

Domain Description

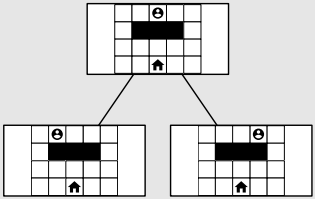

Initial State

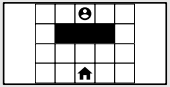

Goal **Hypotheses**

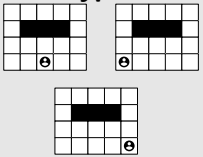

Observations

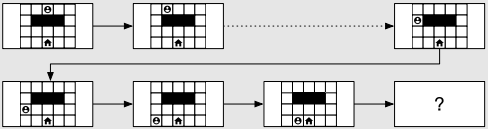

Goal/Plan Recognition problems have **four** key ingredients

Domain Description

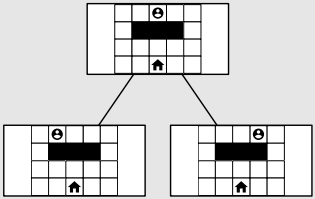

Initial State

Goal **Hypotheses**

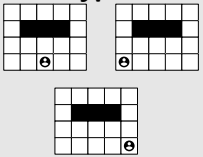

Observations

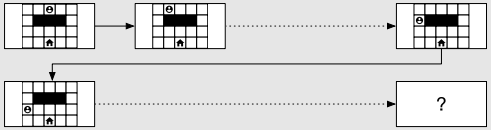

Goal/Plan Recognition problems have **four** key ingredients

Domain Description

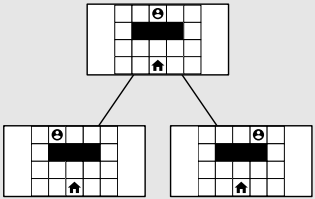

Initial State

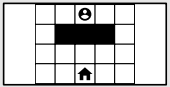

Goal **Hypotheses**

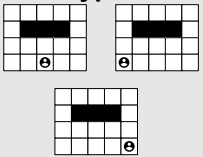

Observations

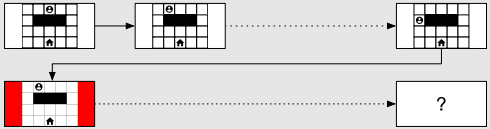

Solution

Correct Goal

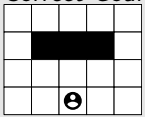

Goal/Plan Recognition problems have **four** key ingredients

Domain Description

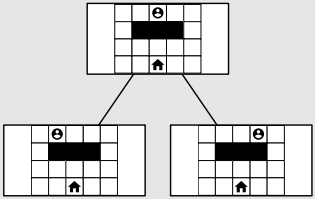

Initial State

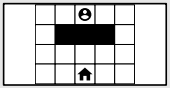

Goal **Hypotheses**

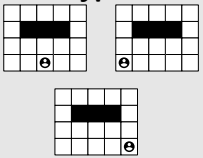

Observations

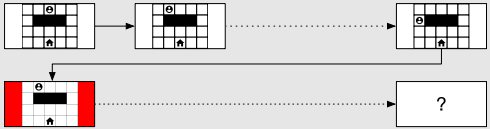

Solution

Probability Distribution

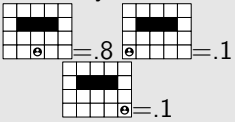

Table of Contents

Generalised Planning in BDI

Overview

- We define a high-level reasoning cycle based on previous work
  - Only declarative goals (no plan library)
  - Generalised planner as the primary means-ends reasoning process
- Key processes:
  - Intention selection
  - Desire filtering

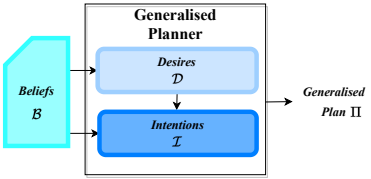

[0] Felipe Meneguzzi and Lavindra de Silva. "Planning in BDI agents: a survey of the integration of planning algorithms and agent reasoning". In: *KER* 30.1 (2015), pp. 1–44.

Definition (Generalised Planning Problem)

A *generalised planning problem* $\mathcal{GP} = \langle \mathcal{P}_0, \mathcal{P}_1, ..., \mathcal{P}_N \rangle$ is a set of planning problems ($N \geq 2$), where each problem $\mathcal{P}_i = \langle s_0, s_g \rangle$ that share some common structure (typically a planning domain Ξ). A solution to a generalised planning problem is a generalised plan $\Pi_{\mathcal{GP}}$, which when followed will solve any problem in $\mathcal{GP}$.

Generalised Planning problems have three key ingredients

Domain Description

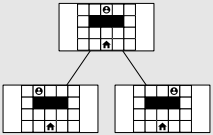

Initial States

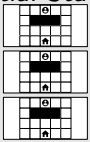

Goal States

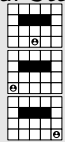

Solution

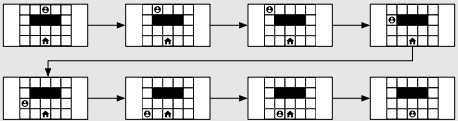

Generalised Planning problems have three key ingredients

Domain Description

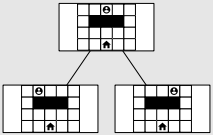

Initial States

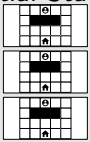

Goal States

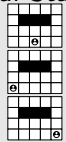

Solution

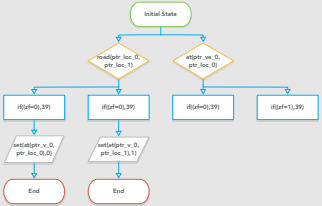

Generalised Planning in BDI

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(π)
7: Find $\{\langle\varphi_1, D_1\rangle \ldots \langle\varphi_n, D_n\rangle\} \in \mathcal{D}^2$
s.t. $\exists \Pi$, $\Pi = \mathcal{G}\text{PLANNER}(\{\langle\Xi, \mathcal{B}, D_1\rangle \ldots \langle\Xi, \mathcal{B}, D_n\rangle\})$
8: $\mathcal{I} \leftarrow \{\langle\langle\varphi_1, D_1\rangle, \Pi\rangle, \langle\langle\varphi_n, D_n\rangle, \Pi\rangle\}$

Generalised Planning in BDI

Key advantages

- Generalised planning problems naturally deal with concurrent desires/intentions
  - Each desire is a sub-problem in $\mathcal{GP}$
  - Resulting generalised plan $\Pi_{\mathcal{GP}}$ analogous to BDI plan-rules
  - Means-ends reasoning inherits properties of the underlying plans
- Allows us to reason about BDI agent behaviour using goal recognition
  - Further automation in multiagent systems

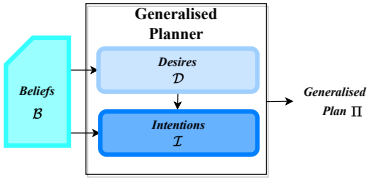

Table of Contents

All together now

Generalised Intent Recognition and BDI as a Theory of Mind

- We define a generalised goal recognition problem $\langle \mathbb{G}, \Omega_\Pi \rangle$, where $\mathbb{G} = \langle \mathcal{GP}_0, \mathcal{GP}_1, ..., \mathcal{GP}_N \rangle$
- Solving this problem consists of computing posterior probabilities over $\mathbb{G}$ given Ω_Π:

$$\mathbb{P}(\mathcal{GP} \mid \Omega_\Pi) = \eta * \mathbb{P}(\Omega_\Pi \mid \mathcal{GP}) * \mathbb{P}(\mathcal{GP})$$

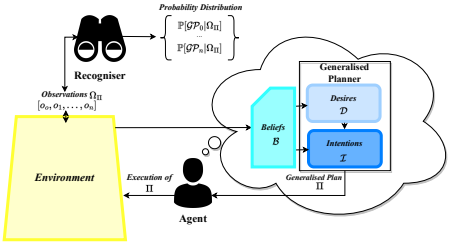

All together now

Generalised Intent Recognition and BDI as a Theory of Mind

- We define a generalised goal recognition problem $\langle \mathbb{G}, \Omega_\Pi \rangle$, where $\mathbb{G} = \langle \mathcal{GP}_0, \mathcal{GP}_1, ..., \mathcal{GP}_N \rangle$
- Solving this problem consists of computing posterior probabilities over $\mathbb{G}$ given Ω_Π:

$$\mathbb{P}(\mathcal{GP} \mid \Omega_\Pi) = \eta * \mathbb{P}(\Omega_\Pi \mid \mathcal{GP}) * \mathbb{P}(\mathcal{GP})$$

- BDI reasoning cycle and goal recognition provide an effective *Theory of Mind*

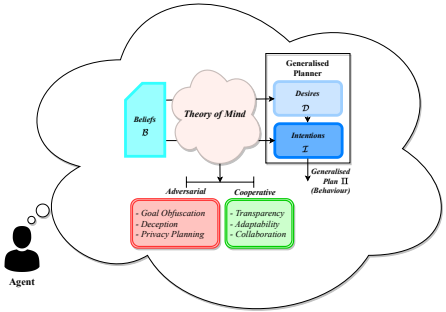

Table of Contents

Promising Approaches

Moving forward

BDI agents using our model now have a model and an inference mechanism to be fully aware of others:

- Adversarial Contexts (counterplanning)
- Cooperative Contexts (transparent planning)

Challenges and Opportunities

- This paper lays out a generic framework, but most of its components are still open research questions
  - Reasoning cycle
    - computationally expensive
    - high-level, no failures, replaning
  - Generalised recognition approaches are still in their infancy
- However, this provides a research agenda for many years to come