Agent-Based Systems

Michael Rovatsos
mrovatso@inf.ed.ac.uk

Lecture 4 – Practical Reasoning Agents

Where are we?

Last time . . .

- Specifying agents in a logical, deductive framework
- General framework, agent-oriented programming, MetateM
- Intelligent autonomous behaviour not only determined by logic!
- (Although this does not mean it cannot be simulated with deductive reasoning methods)
- Need to look for more practical view of agent reasoning

Today . . .

- **Practical Reasoning Systems**

Practical reasoning

- **Practical reasoning** is reasoning directed towards *actions*, i.e. deciding what to do
- Principles of practical reasoning applied to agents largely derive from work of philosopher Michael Bratman (1990):

  > *Practical reasoning is a matter of weighing conflicting considerations for and against competing options, where the relevant considerations are provided by what the agent desires/values/cares about and what the agent believes.*

- Difference to theoretical reasoning, which is concerned with *belief* (e.g. reasoning about a mathematical problem)
- Important: computational aspects (e.g. agent cannot go on deciding indefinitely, he has to act)
- Practical reasoning is foundation for **Belief-Desire-Intention model** of agency

Practical reasoning

- Practical reasoning consists of two main activities:
  1. Deliberation: deciding *what* to do
  2. Means-ends reasoning: deciding *how* to do it
- Combining them appropriately = foundation of deliberative agency
- **Deliberation** is concerned with determining what one wants to achieve (considering preferences, choosing goals, etc.)
- Deliberation generates **intentions** (interface between deliberation and means-ends reasoning)
- **Means-ends reasoning** is used to determine how the goals are to be achieved (thinking about suitable actions, resources and how to "organise" activity)
- Means-ends reasoning generates **plans** which are turned into actions

Intentions

- In ordinary speech, intentions refer to actions or to states of mind; here we consider the latter
- We focus on *future-directed intentions* i.e. *pro-attitudes* that tend to lead to actions
- We make *reasonable attempts* to fulfil intentions once we form them, but they may change if circumstances do
- Main properties of intentions:
  - **Intentions drive means-ends reasoning**: If I adopt an intention I will attempt to achieve it, this affects action choice
  - **Intentions persist**: Once adopted they will not be dropped until achieved, deemed unachievable, or reconsidered
  - **Intentions constrain future deliberation**: Options inconsistent with intentions will not be entertained
  - **Intentions influence beliefs concerning future practical reasoning**: Rationality requires that I believe I can achieve intention

Intentions

- Bratman's model suggests the following properties:
  - Intentions pose problems for agents, who need to determine ways of achieving them
  - Intentions provide a 'filter' for adopting other intentions, which must not conflict
  - Agents track the success of their intentions, and are inclined to try again if their attempts fail
  - Agents believe their intentions are possible
  - Agents do not believe they will not bring about their intentions
  - Under certain circumstances, agents believe they will bring about their intentions
  - Agents need not intend all the expected side effects of their intentions

Intentions

- Cohen-Levesque theory of intentions based on notion of **persistent goal**
- An agent has a persistent goal of φ iff:
  1. It has a goal that φ eventually becomes true, and believes that φ is not currently true
  2. Before it drops the goal φ, one of the following conditions must hold:
     - the agent believes φ has been satisfied
     - the agent believes φ will never be satisfied
- Definition of intention (consistent with Bratman's list):

  An agent intends to do action α iff it has a persistent goal to have brought about a state wherein it believed it was about to do α, and then did α.

Desires

- Desires describe the states of affairs that are considered for achievement, i.e. basic preferences of the agent
- Desires are much weaker than intentions, they are not directly related to activity:

> *My desire to play basketball this afternoon is merely a potential influence of my conduct this afternoon. It must vie with my other relevant desires [. . .] before it is settled what I will do. In contrast, once I intend to play basketball this afternoon, the matter is settled: I normally need not continue to weigh the pros and cons. When the afternoon arrives, I will normally just proceed to execute my intentions. (Bratman, 1990)*

The BDI Architecture

Sub-components of overall BDI control flow:

- Belief revision function
  - Update beliefs with sensory input and previous belief
- Generate options
  - Use beliefs and existing intentions to generate a set of alternatives/options (=desires)
- Filtering function
  - Choose between competing alternatives and commit to their achievement
- Planning function
  - Given current belief and intentions generate a plan for action
- Action generation: iteratively execute actions in plan sequence

The BDI Architecture

Deliberation process in the BDI model:

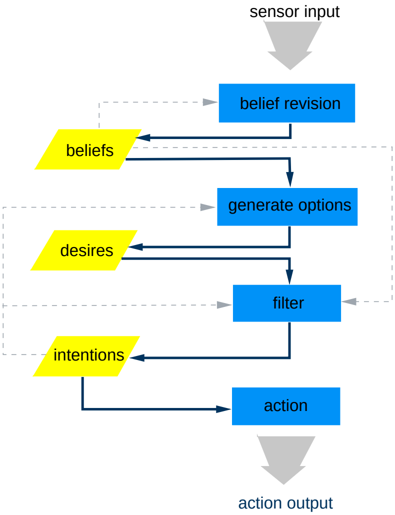

The BDI architecture – formal model

- Let $B \subseteq Bel$, $D \subseteq Des$, $I \subseteq Int$ be sets describing beliefs, desires and intentions of agent
- Percepts *Per* and actions *Ac* as before, *Plan* set of all plans (for now, sequences of actions)
- We describe the model through a set of abstract functions
- Belief revision $brf : \wp(Bel) \times Per \rightarrow \wp(Bel)$
- Option generation $options : \wp(Bel) \times \wp(Int) \rightarrow \wp(Des)$
- Filter to select options $filter : \wp(Bel) \times \wp(Des) \times \wp(Int) \rightarrow \wp(Int)$
- Means-ends reasoning: $plan : \wp(Bel) \times \wp(Int) \times \wp(Ac) \rightarrow Plan$

BDI control loop (first version)

Practical Reasoning Agent Control Loop

$B \leftarrow B_0$; $I \leftarrow I_0$; /* initialisation */
while *true* do
    get next percept ρ through *see*(...) function
    $B \leftarrow brf(B, \rho)$; $D \leftarrow options(B, I)$; $I \leftarrow filter(B, D, I)$;
    $\pi \leftarrow plan(B, I, Ac)$;
    while not ($empty(\pi)$ or $succeeded(I, B)$ or $impossible(I, B)$) do
        $\alpha \leftarrow head(\pi)$;
        $execute(\alpha)$;
        $\pi \leftarrow tail(\pi)$;
    end-while
end-while

Means-ends reasoning

- So far, we have not described *plan* function, i.e. *how* to achieve goals (ends) using available means
- Classical **AI planning** uses the following representations as inputs:
    - A **goal** (intention, task) to be achieved (or maintained)
    - Current **state** of the environment (beliefs)
    - **Actions** available to the agent
- Output is a **plan**, i.e. “a recipe for action” to achieve goal from current state
- STRIPS: most famous classical planning system
    - State and goal are described as logical formulae
    - Action schemata describe preconditions and effects of actions

Blocks world example

- Given: A set of cube-shaped blocks sitting on a table
- Robot arm can move around/stack blocks (one at a time)
- Goal: configuration of stacks of blocks
- Formalisation in STRIPS:
- State description through set of literals, e.g.

$$\{Clear(A), On(A, B), OnTable(B), OnTable(C), Clear(C)\}$$

- Same for goal description, e.g.

$$\{OnTable(A), OnTable(B), OnTable(C)\}$$

- Action schemata: precondition/add/delete list notation

Blocks world example

- Some action schemata examples

$Stack(x, y)$
$pre\ \{Clear(y), Holding(x)\}$
$del\ \{Clear(y), Holding(x)\}$
$add\ \{ArmEmpty, On(x, y)\}$

$UnStack(x, y)$
$pre\ \{On(x, y), Clear(x), ArmEmpty\}$
$del\ \{On(x, y), ArmEmpty\}$
$add\ \{Holding(x), Clear(y)\}$

$Pickup(x)$
$pre\ \{Clear(x), OnTable(x), ArmEmpty\}$
$del\ \{OnTable(x), ArmEmpty\}$
$add\ \{Holding(x)\}$

$PutDown(x)$
$pre\ \{Holding(x)\}$
$del\ \{Holding(x)\}$
$add\ \{ArmEmpty, OnTable(x)\}$

- (Linear) plan = sequence of action schema instances
- Many algorithms, simplest method: state-space search

Formal model of planning

- Define a **descriptor** for an action $\alpha \in Ac$ as

$$\langle P_\alpha, D_\alpha, A_\alpha \rangle$$

  defining sets of first-order logic formulae of precondition, delete- and add-list
- Although these may contain variables and logical connectives we ignore these for now (assume ground atoms)
- A **planning problem** $\langle \Delta, O, \gamma \rangle$ over Ac specifies
  - Δ as the (belief about) initial state (a list of atoms)
  - a set of operator descriptors $O = \{\langle P_\alpha, D_\alpha, A_\alpha \rangle | \alpha \in Ac\}$
  - an intention γ (set of literals) to be achieved
- A **plan** is a sequence of actions $\pi = (\alpha_1, \ldots \alpha_n)$ with $\alpha_i \in Ac$

Formal model of planning

- In a planning problem $\langle \Delta, O, \gamma \rangle$ a plan π determines a sequence of **environment models** $\Delta_0, \ldots, \Delta_n$
- For these, we have
  - $\Delta_0 = \Delta$ and
  - $\Delta_i = (\Delta_{i-1} \backslash D_{\alpha_i}) \cup A_{\alpha_i}$ for $1 \leq i \leq n$
- π is **acceptable** wrt $\langle \Delta, O, \gamma \rangle$ iff $\Delta_{i-1} \models P_{\alpha_i}$ for all $1 \leq i \leq n$
- π is **correct** wrt $\langle \Delta, O, \gamma \rangle$ iff π is acceptable and $\Delta_n \models \gamma$
- The problem of AI planning:

  Find a correct plan π for planning problem $\langle \Delta, O, \gamma \rangle$ if one exists, else announce that none exists

Formal model of planning

- Below, we will use
  - *head*(π), *tail*(π), *pre*(π), *body*(π) to refer to parts of a plan,
  - *execute*(π) to denote execution of whole plan,
  - *sound*(π, *I*, *B*) to denote that π is correct given intentions *I* and beliefs *B*
- Note: planning does not have to involve plan *generation*
- Alternatively, plan libraries can be used
- Now we are ready to integrate means-ends reasoning in our BDI implementation

BDI control loop (first version)

Practical Reasoning Agent Control Loop

$B \leftarrow B_0$; $I \leftarrow I_0$; /* initialisation */
while *true* do
    get next percept ρ through *see*$(\ldots)$ function
    $B \leftarrow brf(B, \rho)$; $D \leftarrow options(B, I)$; $I \leftarrow filter(B, D, I)$;
    $\pi \leftarrow plan(B, I, Ac)$;
    while not ($empty(\pi)$ or $succeeded(I, B)$ or $impossible(I, B)$) do
        $\alpha \leftarrow head(\pi)$;
        $execute(\alpha)$;
        $\pi \leftarrow tail(\pi)$;
    end-while
end-while

Commitment to ends and means

- We should think that deliberation and planning are sufficient to achieve desired behaviour, unfortunately things are more complex
- After filter function, agent makes a **commitment** to chosen option (this implies *temporal persistence*)
- Question: how long should an intention persist? (remember dung beetle?)
- Different commitment strategies:
    - Blind/fanatical commitment: maintain intention until it has been achieved
    - Single-minded commitment: maintain intention until achieved or impossible
    - Open-minded commitment: maintain intention as long as it is believed possible
- Note: agents commit themselves both to ends (intention) and means (plan)

Commitment to ends and means

- As concerns commitment to means, we choose single-minded commitment (using predicates $succeeded(I, B)$ and $impossible(I, B)$)
- Commitment to ends: **intention reconsideration**
  - When whould we stop to think whether intentions are already fulfilled/impossible to achieve?
  - Trade-off: intention reconsideration is costly but necessary
    ➡ meta-level control might be useful ($reconsider(I, B)$ predicate)
- When is an IR strategy optimal (given that planning and intention choice are)?
- IR strategy is optimal if it *would* have changed intentions *had* he deliberated again (this assumes IR itself is cheap . . .)
- Rule of thumb: being “bold” is fine as long as world doesn’t change at a high rate

BDI control loop (second version)

Practical Reasoning Agent Control Loop

$B \leftarrow B_0$; $I \leftarrow I_0$; /* initialisation */
while *true* do
    get next percept ρ through $see(\ldots)$ function
    $B \leftarrow brf(B, \rho)$; $D \leftarrow options(B, I)$; $I \leftarrow filter(B, D, I)$;
    $\pi \leftarrow plan(B, I, Ac)$;
    while not ($empty(\pi)$ or $succeeded(I, B)$ or $impossible(I, B)$) do
        $\alpha \leftarrow head(\pi)$;
        $execute(\alpha)$;
        $\pi \leftarrow tail(\pi)$;
        get next percept ρ though $see(\ldots)$ function
        $B \leftarrow brf(B, \rho)$;
        if $reconsider(I, B)$ then
            $D \leftarrow options(B, I)$; $I \leftarrow filter(B, D, I)$
        if not $sound(\pi, I, B)$ then
            $\pi \leftarrow plan(B, I, Ac)$
    end-while
end-while

Summary

- Discussed practical reasoning systems
- Today the prevailing paradigm in deliberative agency
- Deliberation: an interaction between beliefs, desires and intentions
- Special properties of intentions, C-L theory
- Means-ends reasoning and planning
- Commitment strategies and intention reconsideration
- Next time: **Reactive and Hybrid Agent Architectures**