Autonomous Agents as Intentional Systems

Autonomous Systems
Sistemi Autonomi

Andrea Omicini

`andrea.omicini@unibo.it`

Dipartimento di Informatica – Scienza e Ingegneria (DISI)
Alma Mater Studiorum – Università di Bologna

Academic Year 2017/2018

Next in Line. . .

Intelligent Agents I

According to a classical definition, an intelligent agent is a computational system capable of autonomous action and perception in some environment

Reminder: computational autonomy

- agents are autonomous as they encapsulate (the thread of) control
- control does not pass through agent boundaries
  - only data (knowledge, information) crosses agent boundaries
- agents have no interface, cannot be controlled, nor can they be invoked
- looking at agents, MAS can be conceived as an aggregation of multiple distinct loci of control interacting with each other by exchanging information

Intelligent Agents II

Question: what about the other notions of autonomy?

- autonomy with respect to other agents – social autonomy
- autonomy with respect to environment – interactive autonomy
- autonomy with respect to humans – artificial autonomy
- autonomy with respect to oneself – moral autonomy
- . . .

Question: what is intelligence to autonomy?

- any sort of intelligence?
- which intelligence for which autonomy?
- which intelligent architecture for which autonomy?
- . . .

Actions and Perception

- actions are supposed to change the environment in order to meet agents design objectives
- perception is a process by which the agent recognises the state of the environment, so as to be able to adapt its behaviour to it

[Russell and Norvig, 2002]

- the perception function is the agent's ability to observe its environment
    - the outcome of the perception function is a percept

$$Perception : E \rightarrow Per$$

    which maps environment states to percepts
- the action function represents the agent's decision making process
    - the outcome of the action function is an action

$$Action : Per^* \rightarrow A_c$$

   which maps sequences of percepts to actions

Autonomy

- as we know, the main point about agents is they are *autonomous*: capable of acting independently, encapsulating control over their internal state.
  - thus, an agent is a computer system capable of autonomous actions in some environment
- what about agent behaviour? The most typically-mentioned features are
  - reactivity
  - situatedness
  - pro-activeness
  - social ability

 which, as we already know, can be somehow reduced to autonomy itself

Reactivity

- application in real world domains are characterised by highly dynamic conditions: situations change, information is incomplete, resources are scarce, the actions performed are not deterministic in their effects

- a *reactive* system is one that maintains an ongoing interaction with its environment, and responds to changes that occur in it—in time for the response to be useful

- purely reactive agents decide what to do without reference to their history

- reaction is entirely based on the present, with no reference to past states
    - e.g. *stimolous → response* rules

        $$action : E \rightarrow A_c$$

   - a thermostat is a purely reactive agent

Situatedness

- reactive models and state-less agency are not enough for entities engaged in dynamic environments
- such entities continually face with external events requiring adequate services and behavioural responses
- any "ground" model of action is strictly coupled with the context where the action takes place
- an agent comes with its own model of action
- any agent is then strictly coupled with the environment where it lives and (inter)acts by the very actions it is capable of

Proactiveness

- we need agents able to smartly adapt to their environment
  - autonomously (re)adapting to *changes*
  - recognising *opportunities*
  - *goal-oriented* behaviour
- proactiveness is a *generative* approach
  - agents generate their objectives, and attempt to achieve them
  - agents encapsulate control, and the rule to govern it
  - agents are not driven solely by stimuli
  - agents take the initiative and make something happen, rather than waiting for something to happen

Social Ability

- MAS can be conceived as an articulated world
    - where we cannot go around attempting to achieve goals without taking other entities into account
- some goals can only be achieved with the cooperation of others
- thus, *social ability* in agents is the ability to interact with other agents (and possibly humans) via some kind of agent-communication means
    - speech acts [Searle, 1969]
    - artifact-based interactions [Omicini et al., 2004]
    - signals
    - environment traces [Parunak, 2006]
    - . . .
- coordination is social: cooperation, collaboration, but also competition with others

Agents with State

- to face these growing complexities, stateful agents can be conceived
- these agents have some internal data structure, which is typically used to record information about the environment state and history
- let I be the set of all internal states of the agent: the *perception* function is the same

$$Perception : E \rightarrow Per$$

  while the action-selection function action is defined as a mapping

$$Action : I \rightarrow A_c$$

  from internal states to actions
- an additional function *next* is introduced, which maps an internal state and percept to an internal state:

$$Next : I \times Per \rightarrow I$$

Implementing an Agent with Internal State

Agent control loop (v. 1)

```
1: while true
2:     observe the environment state 'e' and generates a perception(e);
3:     update internal state model:
           i ::= next(i, perception(e));
4:     select an action to execute:
           action(i);
5: end while
```

- we now have the problem to define such agent states
- how to build an effective *internal state model*?
- how to make it run?
- how to update it?

Next in Line. . .

Intentional Systems

- an idea is to refer to human attitudes as *intentional notions*
- when explaining human activity, it is often useful to make statements such as the following:
  - Seb got rain tires because he believed it was going to rain
  - Kimi is working hard because he wants to win world championship again
- these statements can be read in terms of *folk psychology*, by which human behaviour can be explained and can be *predicted* through the attribution of mental attitudes, such as believing and wanting (as in the above examples), hoping, fearing, and so on.

The Intentional Stance

- the philosopher – cognitive scientist – Daniel Dennett coined the term intentional system to describe entities *'whose behaviour can be predicted by the method of attributing to it belief, desires and rational acumen'* [Dennett, 1971]
- Dennett identifies several grades of intentional systems:
  1. a first-order intentional system has beliefs, desires, etc.
     - Seb believes P
  2. a second-order intentional system has beliefs, desires, etc. about beliefs, desires, etc. both its own and of others
     - Seb believes that Kimi believes P
  3. a third-order intentional system is then something like
     - Seb believes that Kimi believes that Seb believes P
  4. . . .

The Intentional Stance in Computing

What entities can be described in terms of *intentional stance*?

- human beings are prone to provide an intentional stance to almost anything
  - sacrifices for ingratiating gods benevolence
  - animism
  - . . .

Ascribing mental qualities to machines [McCarthy, 1979]

Ascribing mental qualities like beliefs, intentions and wants to a machine is sometimes correct if done conservatively and is sometimes necessary to express what is known about its state [...] it is useful when the ascription helps us understand the structure of the machine, its past or future behaviour, or how to repair or improve it

Next in Line. . .

Agents as Intentional Systems

- as computer systems become ever more complex, we need more powerful abstractions and metaphors to explain their operation
- with complexity growing, mechanist / low level explanations become impractical
- the intentional notions can be adopted as abstraction tools, providing us with a convenient and familiar way of describing, explaining, and predicting the behaviour of complex systems
- the idea is to use the intentional stance as an abstraction in computing in order to explain, understand, drive the behaviour—then, crucially, program computer systems

Agents as Intentional Systems

Strong notion of agency

- early agent theorists start from a (strong) notion of agents as intentional systems
- agents were explained in terms of mental attitudes, or mental states
- in their social abilities, agents simplest consistent description implied the intentional stance
- agents contain an explicitly-represented – *symbolic* – model of the world (written somewhere in the working memory)
- agents make decision on what action to take in order to achieve their goals via *symbolic* reasoning

Which Domains for Intention Systems? I

Mental states are a worth *abstraction* for developing agents to effectively act in a class of application *domains* characterised by various practical limitations and *requirements* [Rao and Georgeff, 1995]

Which Domains for Intention Systems? II

- at any instant of time there are many different ways in which an environment may evolve—the environment is not deterministic
- at any instant of time there are many actions or procedures the agent may execute—the agent is not deterministic, too
- at any instant of time the agent may want to achieve several objectives
- the actions or procedures that (best) achieve the various objectives are dependent on the state of the environment—i.e., on the particular situation, context
- the environment can only be sensed locally
- the rate at which computations and actions can be carried out is within *reasonable bounds* to the rate at which the environment evolves

Goal-Oriented & Goal-Directed Systems

- there are two main families of architectures for agents with mental states

  teleo-reactive / goal-oriented agents are based on their own design model and internal control mechanism. The *goal* is not explicitly represented within the internal state, instead it is an 'end state' for agents internal state machine

  deliberative / goal-directed agents are based on symbolic reasoning about goals, which are explicitly represented and processed aside the control loop

Modelling Agents with Mental States I

Modelling agents based on mental states...

- eases the development of agents exhibiting complex behaviour
- provides us with a familiar, non-technical way of understanding and explaining agents
- allows the developer to build MAS by adopting the perspective of a cognitive entity engaged in complex tasks—e.g., what would *I* do in the same situation?
- simplifies the construction, maintenance, and verification of agent-based applications
- is useful when the agent has to comunicate and interact with users or other system entities

Modelling Agents with Mental States II

The intentional stance [Dennett, 2007]

The intentional stance is the strategy of interpreting the behaviour of an entity (person, animal, artifact, whatever) by treating it as if it were a rational agent who governed its 'choice' of 'action' by a 'consideration' of its 'beliefs' and 'desires'.

The scare-quotes around all these terms draw attention to the fact that some of their standard connotations may be set aside in the interests of exploiting their central features: their role in practical reasoning, and hence in the prediction of the behaviour of practical reasoners.

Modelling Agents with Mental States III

Agents with mental states

- agents governing their behaviour based on *internal states* that mimic cognitive (human) *mental states*

  epistemic states representing agents *knowledge*—their knowledge on the world

    - i.e., percepts, beliefs

  motivational states representing agents *objectives*—what they aim to achieve

    - i.e., goals, desires

- the process of selecting one action to execute among the many available based on the actual mental states is called *practical reasoning*

  - i.e., $action(next(i, perception(e))$

Next in Line. . .

Practical vs. Epistemic Reasoning

practical reasoning is reasoning directed towards actions—the process of figuring out what to do in order to achieve what is desired

[Bratman, 1987]

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.

epistemic reasoning is reasoning directed towards knowledge—the process of updating information, replacing old information (no longer consistent with the world state) with new information

Practical Reasoning

- practical reasoning consists of two main cognitive activities

  deliberation when the agent makes decision on what state of affairs the agent desire to achieve

  means-ends reasoning when the agent makes decisions on how to achieve these state of affairs

- the outcome of the deliberation phase are the *intentions*
  - what agent desires to achieve, or what he desires to do
- the outcome of the means-ends reasoning phase is the selection a given course of actions
  - the workflow of the actions the agent intends to adopt in order to achieve its own goals expressed as intentions

Basic Architecture of a Mentalistic Agent
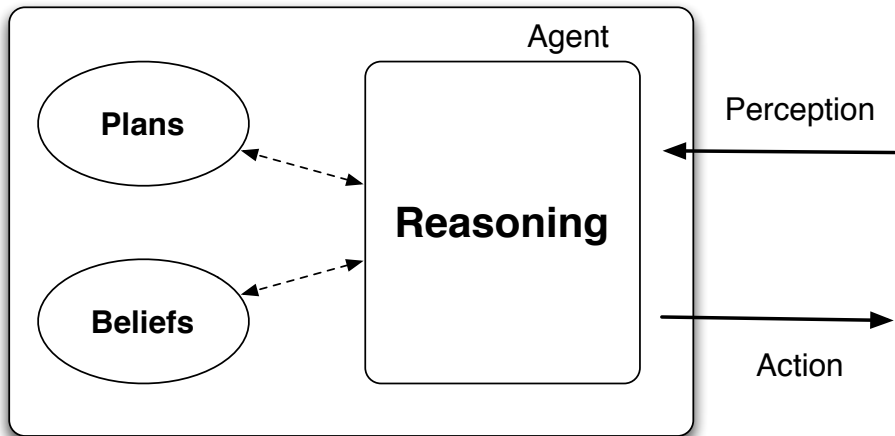
The Role of Intentions in Practical Reasoning I

1. intentions represent a problem to solve for the agent who need to determine how to achieve them
   - if I have an intention to ϕ, you would expect me to *devote resources* to deciding how to bring about ϕ.

2. intentions provide a *filter* for adopting other intentions, which must *not conflict*
   - if I have an intention to ϕ, you would not expect me to adopt an intention ψ such that ϕ and ψ are mutually exclusive

3. intentions tend to be *stable*: agents track the success of their intentions, and are inclined to try again if their attempts fail
   - if an agent's first attempt to achieve ϕ fails, then all other things being equal, it will try an alternative plan to achieve ϕ

4. agents believe their intentions are *possible*
   - that is, they believe that there is at least some way that the intentions could be brought about

The Role of Intentions in Practical Reasoning II

5. agents do *not believe* they will *not bring about* their intentions.
   - it would not be rational for me to adopt an intention to ϕ if I believed ϕ was not possible.

6. under certain circumstances, agents believe they will *bring about their intentions*
   - it would not normally be rational of me to believe that I would bring my intentions about; intentions can fail
   - moreover, it does not make sense that if I believe ϕ is inevitable that I would adopt it as an intention

7. agents need not intend all the expected side effects of their intentions
   - if I believe $\phi \rightarrow \psi$ and I intend that ϕ, I do not necessarily intend ψ also
   - $\rightarrow$ intentions are not closed under implication
   - this last problem is known as the *side effect* or *package deal* problem: I may believe that going to the dentist involves pain, and I may also intend to go to the dentist—but this does not imply in any way that I intend to suffer pain

Intentions vs. Desires

- the adoption of an intention follows the rise of a given *desire*
  - i.e., it follows the adoption of a given goal
- desires and intentions are different concepts
  - "My desire to play basketball this afternoon is merely a potential influencer of my conduct this afternoon. It must live 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]
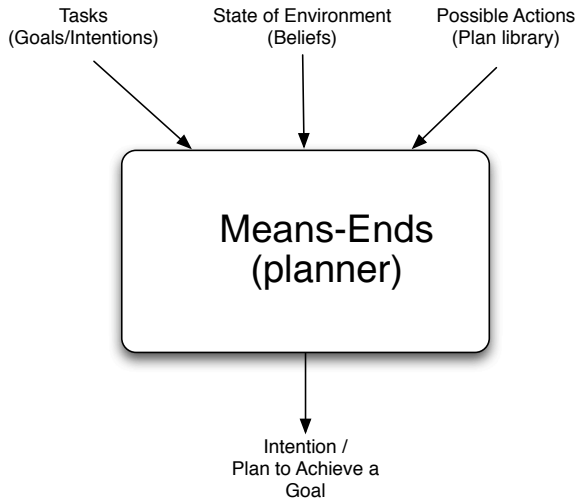
Means-Ends Reasoning I

- the basic idea is to provide agents with three sorts of representations
  - representation of goal / intention to achieve
  - representation of actions / plans – in repertoire
  - representation of the environment

- given the environmental conditions, means-ends reasoning aims at *devising out a plan* that could possibly *achieve the adopted goal* / intention

- the selected intention is an emergent property, reified at runtime by selecting a given plan for achieving a given goal

Means-Ends Reasoning II

Implementing a Practical Reasoning Agent I

Agent control loop (v. 2)

```
1: while true
2:     observe the world;
3:     update internal world model (beliefs);
4:     deliberate which intention to adopt next;
5:     use means-ends reasoning to get a plan for the given intention;
6:     execute the plan;
7: end while;
```

# Implementing a Practical Reasoning Agent: Issues I

Problem

Agents have *bounded resources*—what is called bounded rationality

- deliberation and means-ends processes are not for free: they have *computational costs*
- the time taken to reason and the time taken to act are potentially *unbounded*
- → this harms agent *fitness*—that is, the reactivity and the promptness that is essential for the agent to survive

Implementing a Practical Reasoning Agent: Issues II

- if the agent
    - starts deliberating at t_0
    - begins means-ends at t_1
    - begins executing a plan at t_2
    - ends executing a plan at t_3
- then
    - *time for deliberation* is

 $$t_{deliberation} = t_1 - t_0$$

    - *time for means-ends reasoning* is

    $$t_{meansend} = t_2 - t_1$$

Implementing a Practical Reasoning Agent: Issues III

- *agents environments* are supposed to be highly dynamic
  - many concurrent changes may occur during agent decision-making as well as during the execution of plans
- the deliberated intention is surely worth to be pursued at the precise time when it the deliberation process starts—so, at t_0
- at time t_1, the agent *selects* a goal/intention that would have been optimal if it had been *achieved* at t_0
- the agent runs the risk that the intention selected is no longer optimal – or no longer achievable – by the time the agent has committed to it

Implementing a Practical Reasoning Agent: Issues IV

So, this agent will exhibit an overall optimal behaviour in the following circumstances / under the following conditions:

1. when deliberation and means-ends reasoning take a vanishingly-small amount of time

2. when the world is guaranteed to remain (essentially) static while the agent is deliberating and performing means-ends reasoning, so that the assumptions upon which the choice of intention to achieve and plan to achieve the intention remain valid until the agent has completed both deliberation and means-ends reasoning

3. when an intention that is optimal when achieved at t_0 – the time at which the world is observed – is guaranteed to remain optimal until t_2—the time at which the agent has found a course of action to achieve the intention

Next in Line. . .

The BDI Framework I

According to [Dasgupta and Ghose, 2011]

- one of the most popular and successful framework for agent technology is defined by Rao and Georgeff [Rao and Georgeff, 1992]
- there, the notions of *belief, desire*, and *intention* are the core ones
- hence, agents in this framework are typically referred to as *BDI agents*

The BDI Framework II

beliefs represent at any time the agent's current *knowledge about the world*, including

- information about the current state of the environment inferred from perception devices
- messages from other agents
- internal information

desires represent a *state of the world* the agent is trying *to achieve*

intentions are the chosen means to achieve the agent's desires, and are generally implemented as plans and post-conditions

- as in general it may have *multiple desires*, an agent can have *a number of intentions active at any one time*
- these intentions may be thought of as *running concurrently*, with *one chosen intention active at any one time*

The BDI Framework III

Besides these components, the BDI model includes

plan library — namely, a set of "recipes" representing the *procedural knowledge* of the agent

event queue — where

- *events* — either perceived from the environment or generated by the agent itself to notify an update of its belief base
- *internal subgoals* — generated by the agent itself while trying to achieve a desire

are stored.

The BDI Framework IV

Plans & plan library

- usually, BDI-style agents do no adopt first principles planning at all
- all plans must be *generated* by the agent programmer *at design time*, which are then *selected* for execution at run time
- pre-programmed plans are collected in the plan library
- the planning done by agents consists entirely of context-sensitive subgoal expansion, which is deferred until a subgoal is selected for execution

The BDI Abstract Architecture

Accordingly, the *abstract architecture* proposed by [Rao and Georgeff, 1992] comprise three dynamic and global structures representing agent beliefs, desires, and intentions (BDI), along with an input queue of events

- *update* (write) and *query* (read) operations are possible upon the three structures
    - update operation are subject to compatibility requirements
    - formalised constraints hold upon the mental attitudes
- the events that the system is able to recognise could be either external – i.e., coming from the environment – or internal ones—i.e., coming from some reflexive action
    - events are assumed to be *atomic*, and can be recognised after they have occurred

Implementing a BDI Agent I

Agent control loop (v. 3) [Rao and Georgeff, 1995]

```
1.  initialize-state();
2.  while true do
3.       options := option-generator(event-queue);
4.       selected-options := deliberate(options);
5.       update-intentions(selected-options);
6.       execute();
7.       get-new-external-events();
8.       drop-successful-attitudes();
9.       drop-impossible-attitudes();
10. end-while
```

Implementing a BDI Agent II

1. the agent initialises the internal states

2. the agent enters the main loop

3. the option generator reads the event queue, and returns a list of options

4. the deliberator selects a subset of options to be adopted, and adds these to the intention structure

5. the intentions to be adopted are filtered from the selected ones

6. if there is an intention to perform an atomic action at this point in time the agent executes it

7. any external events that have occurred during the interpreter cycle are then added to the event queue (the same for internal events)

8. the agent modifies the intention and the desire structures by dropping successful ones

9. finally, impossible desires and intentions are dropped, too

Implementing a BDI Agent III

More formally—in some sense

Agent control loop (v. 4)

1. B:=B_0;
2. I:=I_0;
3. *while true do*
4.   *get new percept* Γ;
5.   B:=*brf*(Γ, B);
6.   I:=*deliberate*(B);
8.   π:=*plan*(B, I);
9.   *execute*(π);
10. *end−while*

How does a BDI Agent Deliberate?

Problem

How can we made reasoning procedures of deliberation and option generation sufficiently fast to satisfy the real time demands placed upon the cognitive system?

- deliberation can be decomposed in two phases:

  option generation — understand what are the available alternatives
  deliberation — choose (and filter) between the adoptable goals/intentions

- chosen options are then intentions, so the agents commit to the selected ones—and executes them

Refining Deliberation Function I

option generation — the agent generates a set of possible alternatives; represents option generation via a function, options, which takes agent's current beliefs and current intentions, and from them determines a set of options (i.e., desires)

deliberation — the agent chooses between *competing* alternatives, and commits to the intention to achieving them; in order to select between competing options, an agent uses a *filter* function

Refining Deliberation Function II

Notes

- the strategy for deliberating between goals typically is in the hands of the agent developer
- most BDI programming platforms provide mechanisms to describe under which conditions some goal should inhibit the others (*goal formulae*)
- typically, such goal formulae are first-order logic predicates indicating contexts and trigger conditions
- *game theory* can enter the picture, here: i.e., maximising expected utilities

BDI Agent Control Loop

Agent control loop (v. 5)

1. B:=B_0;
2. I:=I_0;
3. *while true do*
4.   *get new percept* Γ;
5.   B:=*brf*(Γ, B);
6.   D:=*options*(B, I);
7.   I:=*filter*(B, D, I);
8. π:=*plan*(B, I);
9.   *execute*(π);
10. *end−while*

Structure of BDI Systems

BDI architectures are based on the following constructs

- a set of beliefs
- a set of desires (or goals)
- a set of intentions
  - or better, a subset of the goals with an associated stack of plans for achieving them; these are the intended actions
- a set of internal events
  - elicited by a belief change (i.e., updates, addition, deletion) or by goal events (i.e. a goal achievement, or a new goal adoption)
- a set of external events
  - perceptive events coming form the interaction with external entities (i.e. message arrival, signals, etc.)
- a plan library (repertoire of actions) as a further (static) component

Basic Architecture of a BDI Agent [Wooldridge, 2002]
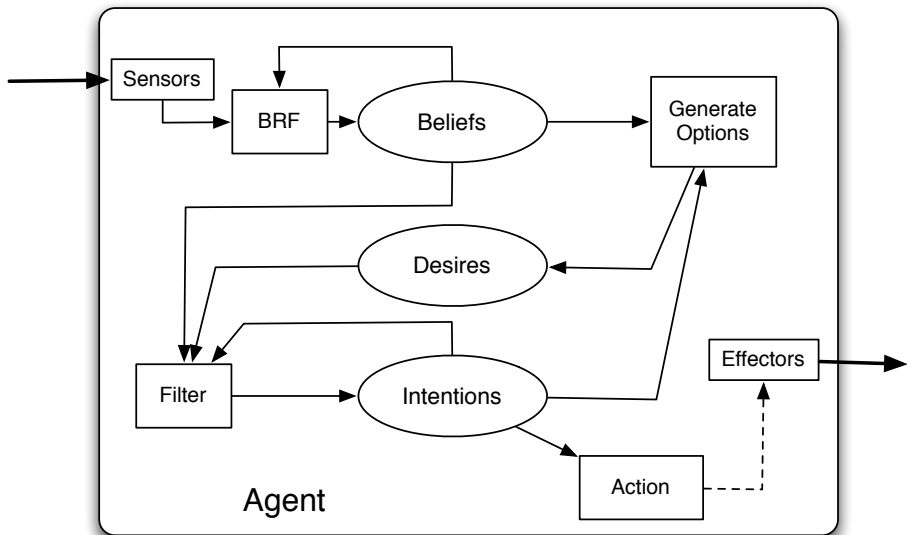
Post-Declarative Systems

It was said that this approach leads to a kind of *post-declarative* programming

- in procedural programming, we say exactly what a system should do
- in declarative programming, we state something that we want to achieve, give the system general info about the relationships between objects, and let a built-in control mechanism (e.g., goal-directed theorem proving) figure out what to do
- with intentional agents, we give a very abstract specification of the system, and let the control mechanism figure out what to do, knowing that it will act in accordance with the built-in theory of agency

Actually, the BDI framework combines in an excellent way both (post)declarative structure and procedural knowledge in terms of plans

Beliefs

Beliefs

Agents knowledge is structured in beliefs about the current state of the world

- they are informational units, typically implemented as ground sets of literals, possibly with no disjunctions or implications
- they should reflect only the information which is currently held (i.e. situated)
- they are expected to change in the future, i.e., as well as the environment changes

Plans

Plans

Plans represent the means the agent has to change the world, and to bring it closer to his desires

- they are language constructs, typically implemented in the form of procedural structures
- plans have a 'body', describing the workflow of activities (actions) that have to be executed for plan execution to be successful
- the conditions under which a plan can be chosen as an option are specified in an *invocation condition* (triggering event) and a *pre- or context- condition* (situation that must hold for the plan to be executable)

Intentions

Intentions

Intentions are emergent properties reified at runtime by selecting a given plan for achieving a given goal

- represented 'on-line' using a run-time stack of hierarchically plans related to the ongoing adopted goals
- similarly to how Prolog interpreter handle clauses
- multiple intention stacks can coexist, either running in parallel, suspended until some condition occurs, or ordered for execution in some way

BDI Viewpoints I

There are three main viewpoints over the BDI Model [Mazal et al., 2008]:

philosophical based on the work of philosopher Bratman [Bratman, 1987], using uses terms of folk psychology to view humans as planning agents: the main concepts in his work are *beliefs* (what an agent knows about the world), *desires* (what the agent wants, can be contradictory) and *intentions* (desires that the agent has decided to reach, cannot be contradictory)

logical mainly Rao and Georgeff's BDI CTL [Rao and Georgeff, 1998] – multimodal logics with possible world semantics –, providing beliefs, goals (desires), and intentions with a precise logical semantics

BDI Viewpoints II

implementation there are a huge number of systems and technologies that are said to conform to the BDI model—between the BDI CTL logics (very expressive) and the implementing systems. which then treat the main modalities rather as data structures, and mostly focus on plans

BDI Agents Programming Platforms

Jason (Brasil) `http://jason.sourceforge.net/`
Agent platform and language for BDI agents based on AgentSpeak(L)

JADEX (Germany) `http://www.activecomponents.org/`
Agent platform for BDI and Goal-Directed Agents

2APL (Netherlands) `http://www.cs.uu.nl/2apl/`
Agent platform providing programming constructs to implement cognitive agents based on the BDI architecture

3APL (Netherlands) `http://www.cs.uu.nl/3apl/`
A programming language for implementing cognitive agents

PRACTIONIST (Italy) `http://practionist.eng.it/`
Framework built on the Bratman's theory of practical reasoning to support the development of BDI agents

ASTRA `http://astralanguage.com/`
A distributed / concurrent programming language based on agent-oriented programming

References I

Bratman, M. E. (1987).
*Intention, Plans, and Practical Reason*.
Harvard University Press.

Bratman, M. E. (1990).
What is intention?
In Cohen, P. R., Morgan, J. L., and Pollack, M. E., editors, *Intentions in Communication*, pages 15–32. The MIT Press, Cambridge, MA.

Dasgupta, A. and Ghose, A. K. (2011).
BDI agents with objectives and preferences.
In Omicini, A., Sardina, S., and Vasconcelos, W., editors, *Declarative Agent Languages and Technologies VIII*, volume 6619 of *Lecture Notes in Computer Science*, pages 22–39. Springer Berlin Heidelberg.

Dennett, D. (1971).
Intentional systems.
Journal of Philosophy, 68:87–106.

References II

Dennett, D. (2007).
Intentional systems theory.
In Beckermann, A. and Walter, S., editors, *Oxford Handbook of the Philosophy of Mind*, chapter 19. Oxford University Press.

Mazal, Z., Kočí, R., Janoušek, V., and Zbořil, F. (2008).
PNagent: A framework for modelling BDI agents using object oriented Petri nets.
In *2008 8th International Conference on Intelligent Systems Design and Applications (ISDA 2008)*, volume 2, pages 420–425. IEEE.

McCarthy, J. (1979).
Ascribing mental qualities to machines.
In Ringle, M., editor, *Philosophical Perspectives in Artificial Intelligence*, Harvester Studies in Cognitive Science, pages 161–195. Harvester Press, Brighton.

Omicini, A., Ricci, A., Viroli, M., Castelfranchi, C., and Tummolini, L. (2004).
Coordination artifacts: Environment-based coordination for intelligent agents.
In Jennings, N. R., Sierra, C., Sonenberg, L., and Tambe, M., editors, *3rd international Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004)*, volume 1, pages 286–293, New York, USA. ACM.

References III

Parunak, H. V. D. (2006).
A survey of environments and mechanisms for human-human stigmergy.
In Weyns, D., Parunak, H. V. D., and Michel, F., editors, *Environments for Multi-Agent Systems II*, volume 3830 of *LNCS*, pages 163–186. Springer.

Rao, A. S. and Georgeff, M. P. (1992).
An abstract architecture for rational agents.
In Nebel, B., Rich, C., and Swartout, W. R., editors, *3rd International Conference on Principles of Knowledge Representation and Reasoning (KR '92)*, pages 439–449, Cambridge, MA, USA. Morgan Kaufmann.
Proceedings.

Rao, A. S. and Georgeff, M. P. (1995).
BDI agents: From theory to practice.
In Lesser, V. R. and Gasser, L., editors, *1st International Conference on Multi Agent Systems (ICMAS 1995)*, pages 312–319, San Francisco, CA, USA. The MIT Press.

Rao, A. S. and Georgeff, M. P. (1998).
Decision procedures for BDI logics.
Journal of Logic and Computation, 8(3):293–342.

References IV

Russell, S. J. and Norvig, P. (2002).
Artificial Intelligence: A Modern Approach.
Prentice Hall / Pearson Education International, Englewood Cliffs, NJ, USA, 2nd edition.

Searle, J. (1969).
Speech Acts: An Essay in the Philosophy of Language.
Cambridge University Press.

Wooldridge, M. J. (2002).
An Introduction to MultiAgent Systems.
John Wiley & Sons Ltd., Chichester, UK.

Autonomous Agents as Intentional Systems

Autonomous Systems
Sistemi Autonomi

Andrea Omicini

andrea.omicini@unibo.it

Dipartimento di Informatica – Scienza e Ingegneria (DISI)
Alma Mater Studiorum – Università di Bologna

Academic Year 2017/2018