The Knowledge Engineering Review, Vol. 31:3, 207–238. © Cambridge University Press, 2016
doi:10.1017/S0269888916000096

BDI agents in social simulations: a survey

CAROLE ADAM[1] and BENOIT GAUDOU[2]

[1]*Grenoble Informatics Laboratory, University Grenoble Alpes IMAG, 700 avenue Centrale, campus universitaire, 38401 Saint-Martin d'Hères France;*
e-mail: carole.adam@imag.fr;
[2]*Toulouse Institute of Computer Science Research, University Toulouse 1 Capitole 2 rue du Doyen Gabriel Marty, 31042 Toulouse Cedex 9, France;*
e-mail: benoit.gaudou@ut-capitole.fr

Abstract

Modelling and simulation have long been dominated by equation-based approaches, until the recent advent of agent-based approaches. To curb the resulting complexity of models, Axelrod promoted the KISS principle: 'Keep It Simple, Stupid'. But the community is divided and a new principle appeared: KIDS, 'Keep It Descriptive, Stupid'. Richer models were thus developed for a variety of phenomena, while agent cognition still tends to be modelled with simple reactive particle-like agents. This is not always appropriate, in particular in the social sciences trying to account for the complexity of human behaviour. One solution is to model humans as belief, desire and intention (BDI) agents, an expressive paradigm using concepts from folk psychology, making it easier for modellers and users to understand the simulation. This paper provides a methodological guide to the use of BDI agents in social simulations, and an overview of existing methodologies and tools for using them.

1 Introduction

After years dominated by equations, that describe the behaviour of the studied system or phenomenon only from a global point of view, modelling and simulation have undergone a deep revolution with the application of multi-agent systems (MAS) to the problem of the modelling and simulation of complex systems. Agent-based modelling and simulation (ABMS) is a recent paradigm that allows modellers to reason and represent the phenomenon at the individual level, and to take into account heterogeneity and complexity both in the individual layer and in the environment layer. ABMS has been successfully used, for example, in Ecology (DeAngelis & Gross, 1992), Economics (Tesfatsion, 2002) or Social Sciences (Gilbert & Troitzsch, 2005).

In the context of agent-based models aimed to 'enrich our understanding of fundamental processes', Axelrod (1997) advocates the KISS principle: 'Keep It Simple, Stupid'. His reasoning is that 'both the researcher and the audience have limited cognitive ability', which limits their ability to understand the model content, and thus their level of confidence in the emergence of a 'surprising result'. That is, by keeping the assumptions underlying the model simple, we can better understand the complexity which arises from the simulation. This notion of simplicity manifests itself in a model as a small number of parameters, a simple environment, a small number of entities and very simple (reactive) agent behaviour. But whether simplicity is *always* appropriate (and if not, when not) has come into question in the multi-agent community, see for example, Drogoul *et al*. (2002). Indeed it has been suggested that there is much to be gained from the simulation of richer, more realistic models; one perspective on this is captured by the KIDS principle (Keep It Descriptive, Stupid).

This principle introduced in Edmonds and Moss (2005) recommends using a formalism as descriptive and straightforward as needed (and possible) when modelling, even if it implies greater complexity.

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

The idea is that it is better to start from a well-adapted formalism, descriptive enough for the needs of the simulation, and to possibly simplify it later, rather than to start from a simple but unadapted formalism that then needs to be extended in an *ad hoc* way to meet the simulation needs. As an example of the KIDS approach, a lot of work has been done on environmental management (with geographic information systems, e.g. Thierry *et al*., 2014) with a large number of parameters, heterogeneity among entities, precise time management and multi-scale simulations.

Nevertheless, the behavioural component usually remains very simple, in particular in ecological or epidemiological simulations (Laperriere *et al*., 2009; Simoes, 2012) that are interested in very simple entities (viruses, bacterias, ducks, etc.), which is their main reason for modelling them with simple and reactive agents. Sun says that

> *a significant shortcoming of current computational social simulation is that most of the work assumes very rudimentary cognition on the part of agents. […] It not only limits the realism, and hence the applicability, of social simulations, but also precludes the possibility of tackling the question of the micro-macro link in terms of cognitive-social interaction (2007: 1, 2).*

Simulations would therefore benefit from more complex agents models. However, Edmonds and Moss (2005) admit that the suitability of the KIDS principle over KISS, the choice between descriptivity and simplicity, highly depends on the context, the goal of the simulation and various parameters. So finally, the KISS approach favouring simplicity over descriptivity is perfectly adapted in some types of simulations, while in other cases the simulation could greatly benefit from using a more descriptive agent model.

In this paper, we will focus on applications of ABMS to the field of Social Sciences, which is interested in more complex entities (namely human beings) in their complex social behaviour and in their interactions with each other and their environment. In that particular field of application, we want to discuss the suitability of more descriptive agent models to represent humans. Indeed, reactive particle-like agents in an electromagnetic field may be adapted in some contexts (e.g. to model crowds); however, modellers tend to use them even in contexts where they are less appropriate than other alternatives. The advantages of more complex and/or descriptive models include keeping the programming intuitive, allowing agents to adapt to the environment, and having agents which can explain their behaviour.

There are several kinds of complex agents that can be used; in this paper we focus in particular on belief, desire and intention (BDI) agents because they offer a straightforward formalisation of the reasoning of human agents with intuitive concepts (beliefs, desires and intentions) that closely match how humans usually explain their reasoning. For example, Norling argues that

> *'BDI agents have been used with considerable success to model humans and create human-like characters in simulated environments. A key reason for this success is that the BDI paradigm is based in folk psychology, which means that the core concepts of the agent framework map easily to the language people use to describe their reasoning and actions in everyday conversations'. Such a straightforward matching between the concepts being modelled and the formal concepts used in the model is exactly what is advocated by the KIDS principle. It also provides numerous advantages: 'the close match between the core BDI concepts and the natural terminology that the experts being modelled used to describe their reasoning facilitated knowledge capture, representation and debugging' (2004: 1).*

Despite these arguments in favour of using BDI agents in social simulations, some features also deter from using them, in particular their complexity.

To summarise, in this paper we argue in favour of the KIDS approach in a particular context which makes it intuitively more adapted: social science simulations which are interested in human behaviour. In this field, the KIDS approach recommends the use of more complex and descriptive agent models. We focus in particular on BDI agents, since they use folk psychology concepts that straightforwardly match human reasoning as people understand it, making the models easier to design and to understand. However, BDI agents are complex, heavy and not always justified. It is therefore necessary to precisely identify the types of simulations where modellers should take advantage of the benefits they provide, and the other types of simulations where the inconveniences outweigh these benefits. We start with further

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

motivating our focus on BDI agents and defining their features and benefits for social sciences simulations (Section 2). We then go on to provide a methodological guide for modellers regarding the use of BDI agents in their simulation (Section 3). We discuss the literature and compare existing simulations on different axis: the required characteristics in the agents (e.g. learning, emotions, Section 3.1); the field of application (Section 3.2); the goal of the simulation (e.g. training, knowledge discovery, Section 3.3); or the desired level of observation (Section 3.4). We also discuss various features that could deter from using BDI agents in simulations (Section 3.5). Finally, Section 4 presents a short overview of available methodologies and tools for concretely integrating BDI agents into simulations.

2 Motivations for BDI agents in social simulations

Agent-based models employ a great variety of agent architectures, from the simplest particle architectures to the most complex cognitive architectures. In this section we illustrate this variety with a few examples. We then justify our choice of the BDI architecture over other types of 'intelligent' agent architectures, and finally give more details about BDI agents and their benefits.

2.1 Examples of agent-based modelling and simulation with various kinds of agents

2.1.1 Simulation of escaping panic

Helbing *et al*. (2000) propose a model of crowd behaviour in the case of emergency escape. They aim to model panic situations where crowd pressure can be dangerous and to find solutions for the evacuation of dangerous places (e.g. smoke-filled rooms). The authors chose to represent the agents as particles in a force field: all physical and social interactions with other agents and the environment are represented as forces. The movement of particles representing the human beings are described by electromagnetic-like equations.

2.1.2 Simulation of pedestrian mobility after an earthquake

Van Truong *et al*. (2013) propose a model of urban pedestrian mobility after an earthquake in Lebanon, in order to show the importance of population preparedness to reduce casualties. In a data-based geographical space composed of streets, buildings and green spaces, agents representing human beings behave in order to evacuate to the green spaces. Their behaviour is driven by a set of rules influenced by agent features (gender, age, etc.) and by stimuli from the environment and from the other agents.

2.1.3 Simulation of survival in a group

Cecconi and Parisi (1998) propose the use of simulation to evaluate various survival strategies of individuals in a social group. Surviving (and reproducing) or dying depends on the resources available to the agent. These resources can be physical resources (such as food, money), other agents (such as a reproduction partner) or personal skills (such as fighting or working capability). Agents can use two types of strategies: individual survival strategies and social survival strategies; a particularly interesting social strategy consists in giving part of what the agent collects to supply a common pool of resources. The authors were interested not only in the evolution of skills over generations but also in the ideal size of the social groups. Each agent is represented as a neural network. The evolution is modelled by a genetic algorithm aimed at improving the capability to find and produce food.

2.1.4 Crowd simulation for emergency response

Shendarkar *et al*. (2008) model a crowd evacuation under terrorist bomb attacks in a public area. They investigate the influence of various important parameters (e.g. density or size of the crowd, number of policemen driving the evacuation, number of exits in the room) on the quality of the evacuation (e.g. evacuation time, number of casualties). Agents are implemented with an extended BDI architecture, which includes an emotional component and a real-time planner. The behaviour of the agents is captured from a human evolving in a Virtual Reality laboratory (participative modelling).

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

2.1.5 *Revamping the simulation of the survival in a group*

Sun and Naveh (2007) propose to give genuine cognitive capabilities to the agents involved in Cecconi's simulation (Cecconi & Parisi, 1998), in order to make it more cognitively realistic. Moreover, they introduce social institutions that impose obligations on the agents and sanctions in case of violation; for example, an agent can be punished if it refuses to contribute food to the central store when required. Through these additions to the original model, the authors aim at studying the interactions between individual cognition and social institutions. The cognitive processes of the agents are implemented by using the CLARION cognitive architecture, mainly based on Sun *et al*. (2001) work in cognitive sciences. It implements mechanisms for implicit vs. explicit learning, among numerous other cognitive capabilities.

2.1.6 *Summary*

These few examples show a huge variety in the agent architectures. Indeed, they illustrate respectively: a particle architecture; a rule-based architecture; a neural network architecture; a BDI architecture; and a cognitive architecture. We can make a dichotomy between these five architectures *w.r.t.* the concepts manipulated to implement the behaviours and their complexity level.

The first three architectures (particle, rule-based and neural network) do not involve any cognitive concepts in the reasoning and decision-making process. Even if agents represent human beings, their behaviours are for the most part a function of the external stimuli (e.g. the external forces in the first example, the perception of the environment in the second and third examples). These architectures are traditionally called *reactive agents* (Wooldridge, 2009). They are very well adapted to represent simple entities without complex reasoning capabilities (e.g. ants) or to model human beings when focussing on simple behaviour, but they cannot deal with complex reasoning mechanisms involved in more complex human behaviour[1].

In contrast, the last two architectures represent two kinds of *intelligent agents* (Wooldridge, 2009): BDI agents and agents with cognitive architectures. The expression 'cognitive architecture' used here refers to architectures based on cognitive theories, such as SOAR, ACT-R or CLARION (Anderson, 1983; Newell, 1990; Sun, 2002); they are often very complex and based on low-level physiological or neurological concepts. It should not be confused with the term 'cognitive agents' that is usually synonym to 'intelligent agents' in the MAS literature. In this wider sense, agents with a BDI architecture are also considered cognitive; but the BDI architecture provides a description of behaviour at a much higher level, in terms of mental attitudes. Intelligent agent architectures allow agents to be more autonomous and proactive, not only reacting to environmental stimuli, therefore matching humans more closely. They enable the modelling of more complex issues than would be possible with purely reactive agents (as shown by the progression between Sections 2.1.3 and 2.1.5), making them very relevant to our aim of integrating much more descriptive agents in social simulations. Finally, their strong theoretical grounding can provide many guidelines in the modelling of agents, but also many constraints.

2.2 *Choice of the agent architecture*

2.2.1 *Another survey of agent architectures*

Balke and Gilbert (2014) recently published a survey of 14 decision-making architectures for agents, from very simple to very complex, aiming at providing guidelines regarding which one of them should be used for a given research question. One of the architectures they study is the BDI one, some others are various extensions of the standard BDI (with emotions, norms, etc.).

On the contrary in this paper, we focus only on the BDI architecture. We started from the observation that the agents used in simulations are too simple, and more descriptive agents would bring many benefits in realism and validity of results. In the next paragraph we explain why we focus on BDI rather than cognitive architectures such as SOAR or ACT-R.

Our goal is therefore different in that we specifically argue for the wider use of BDI agents in social simulations. We also go beyond Balke and Gilbert's review by providing a technological guide showing

[1] Please note that neural networks can also be used to generate proactive behaviour, for instance Murata *et al*. (2014).

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

interested designers not only why they should use BDI agents, but also how to use them. Indeed, even when designers are aware of the benefits that BDI agents could bring in their simulation, they are often limited by technical difficulties, in particular a lack of methodologies and tools.

2.2.2 BDI vs. cognitive architectures

Both BDI architectures and cognitive architectures allow modellers to deal with the cognitive processes of the agents and to provide the descriptivity we (after Edmonds & Moss, 2005) claim is needed in some social simulations. However, these two kinds of architectures are very different.

First of all, they differ on their conceptual inspiration, that is, the origin of the underlying concepts. Cognitive architectures are based on cognitive sciences and aim at describing human cognitive processes as precisely as possible. This requires in-depth analysis of the functioning of the human brain from a physiological or neurological point of view (e.g. to determine the area of the brain that performs a given function). In contrast, the BDI architecture is based on a philosophical stream. Bratman's (1987) work is formalised mainly with modal logics (Cohen & Levesque, 1990; Rao & Georgeff, 1991). These logics describe the three main intentional states (beliefs, desires, intentions) that lead the agent from its world knowledge, to the choice of the action to perform, and to the execution of these actions. The BDI architecture therefore offers a more straightforward description making models easier to understand for modellers and end-users.

As a result, BDI and cognitive architectures also differ on the abstraction level of the concepts manipulated. BDI concepts are higher-level concepts (mental attitudes), that actually result from the low-level concepts and principles described in cognitive architectures (e.g. memory chunks in ACT-R). BDI architectures can thus be viewed as the explicit and symbolic part of cognitive architectures. BDI architectures focus on the conscious level of one's mind, at a level where human beings can consciously manipulate the concepts and communicate them to others (Norling, 2004). This allows the easy mapping of knowledge from experts to agent behaviour, and helps non-computer scientists to understand agent behaviour.

Finally, they differ on their degree of theoretical standardisation. Each cognitive architecture implementation builds upon its own theory: for example, CLARION (Sun *et al.*, 2001) is based on Sun's work (2002); SOAR (Rosenbloom *et al.*, 1993) is based on Newell's work (Newell, 1990); and ACT-R (Anderson & Lebiere, 1998) is based on Anderson's work (Anderson, 1983, 1993). They have no common focus: CLARION is interested in the dichotomy implicit vs. explicit, in particular for the learning process, whereas ACT-R focusses on the activation process. There is no consensus either on the memory component of the architecture, be it on the number of memories or their nature. On the contrary, in all of the substantial number of existing implementations of the BDI framework and formalisations of the underlying modal logic, the BDI trinity is a constant, allowing a huge community of researchers to work on its improvement and extensions.

Sun has argued that cognitive architectures can enhance simulations (Sun, 2006), and proved his point by revamping an existing simulation with a cognitive architecture instead of the initial reactive agents (Sun, 2007). We argue that BDI architectures can similarly enhance simulations, and also provide additional benefits as shown by Norling: 'the close match between the core BDI concepts and the natural terminology that the experts being modelled used to describe their reasoning facilitated knowledge capture, representation and debugging' (2004).

Therefore BDI architectures appear to be both descriptive enough to describe cognitive processes influencing behaviour, and intuitive enough to be easily understood by modellers and non-computer scientists in general. For these reasons, in this paper we are more particularly interested in BDI agents, and want to provide a guide as to when they can and should be used in social sciences simulations. A similar guide could be proposed about when to use cognitive architectures, but this is not the scope of this article.

2.3 Description of a BDI agent

2.3.1 BDI Theory

The BDI model is theoretically grounded in the philosophical work of Bratman (1987) and Dennett (1989) where the three basic components (beliefs, desires and intentions) are defined. The model attempts to

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

capture the common understanding of how humans reason through: *beliefs* which represent the individual's knowledge about the environment and about their own internal state; *desires* or more specifically *goals* (non-conflicting desires which the individual has decided they want to achieve); and *intentions* which are the set of plans or sequence of actions which the individual intends to follow in order to achieve their goals.

Cohen and Levesque (1990) distinguish BDI from purely reactive models by adding a level of commitment to the set of intentions to achieve a long-term goal. This notion of commitment is therefore a critical part of the BDI model. Rao and Georgeff (1991) show there must also be some form of *rational process* for deciding which intentions to select given a set of circumstances. While the specific way in which this is done varies among implementations, it is generally guided by folk psychology principles (Norling, 2004).

So the core functionalities that BDI systems should implement are as follows: a representation of beliefs, desires and intentions; a rational and logical process for selecting intentions; and a flexible and adaptable commitment to the current set of intentions.

2.3.2 *Implemented systems*

There are a substantial number of implemented BDI systems (e.g. JACK (Busetta *et al*., 1999), Jason (Bordini *et al*., 2007), 3APL (Hindriks *et al*., 1999), Gorite (Rönnquist, 2008), JADEX (Pokahr *et al*., 2005), etc.) supporting the creation of agents defined using these three constructs, and several systems extending these to provide additional human reasoning capabilities. Common to all these implementations is the concept of a *plan library* which is a collection of plans that the agent is capable of performing. Each of these *plans* consists of a body, a goal which the plan is capable of achieving, and a context in which the plan is applicable. The plan body can consist of both: actions which can affect the environment directly, and subgoals which will be expanded into other plans when needed. This nesting of goals within plans allows for a plan–goal hierarchy structure to be constructed, with branches from a goal node leading to plans capable of achieving this goal, and branches from a plan node leading to goals which must be achieved to complete this plan.

Implemented BDI systems also contain a BDI execution engine, which guides the agent reasoning process. There are some variations in the algorithms used by different engines, but a basic version is outlined below:

1. *Perception*: find any new events which may have been triggered, either within the agent or externally from the environment.
2. *Update*: update beliefs with new information provided by this event.
3. *Intention revision*: if the change in beliefs means that an intention or goal is no longer valid (either because it has already been achieved or is no longer possible) then remove it from the set, which may result in some form of failure recovery by attempting to achieve it with a different plan.
4. *Plan filtering*: if the intention was revised, determine the set of applicable plans for the current event, which are appropriate given the current context, and are defined as capable of achieving the goal associated with the event.
5. *Plan selection*: if there is no current plan, select a new plan from this set (the selection method varies from selecting the first available plan to selecting the plan with the greatest expected utility) and add the components of its body to the intention list.
6. *Action*: execute the next step of the current plan, which may involve executing an action or expanding a subgoal by triggering a new event.

2.3.3 *Core features and benefits*

The BDI architecture therefore offers the following desirable core features:

- *Adaptability*: plans are not fully specified and can consist of nested subgoals, which allows flexibility in the agent, making it able to adapt to a changing environment.

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

- *Robustness*: the plan hierarchy structure means that if a plan fails, potentially because of environment changes during its execution, the agent is able to recover by executing another applicable plan if one is available.
- *Abstract programming*: the programmer specifies beliefs, desires and intentions at a high level, making it possible to create complex systems while maintaining transparent, easily understandable code.
- *Explainability*: using a goal-based approach as opposed to a task-based approach means that an agent knows why it is doing a specific task. This allows agents to explain their behaviour to the modeller or the user in an intuitive manner.

2.3.4 BDI extensions

Beyond these core features, the BDI framework can be extended to provide additional desired features, and many works were dedicated to such extensions.

For instance, Singh *et al*. (2011) extend BDI with learning, in order to improve plan selection by dynamically adjusting a confidence measure. There has also been work for integrating real-time planners into the BDI architecture to support the creation of plans at runtime using First Principles (de Silva *et al*., 2009) among other techniques. Nair *et al*. (2005) use a BDI model of coordination, combined with POMDP to handle uncertainty, and prove that their agents perform better teamwork.

A large panel of social concepts have been added to the basic BDI framework. Dignum *et al*. (2000) introduce the concept of norms to 'assist in standardising the behaviour of individuals, making it easier to cooperate and/or interact within that society'. In the same way human behaviour is very much influenced by social norms, agent behaviour can be influenced by norms too, allowing agents to have a better understanding of what other agents are likely to do. Gaudou (2008) focusses on the introduction of social attitudes (and in particular group belief), and (Lorini *et al*., 2009) go one step further by modelling the group acceptance in an institutional context. Trust has also been introduced in several BDI works: Taibi (2010) has implemented trust in a BDI interpreter to help agents to delegate tasks; while Koster *et al*. (2012) allow the agent to reason about its trust model and to adapt it to a given context.

Adam *et al*. extend the BDI framework with emotions by formalising their triggering conditions (Adam *et al*., 2006, 2009), and propose an associated reasoning engine for BDI agents (Adam & Lorini, 2014). Other works had the same aim: Pereira *et al*. (2005) present an extension to the BDI architecture to include Artificial Emotions affecting an agent's reasoning in a similar way to human emotions; and Jones *et al*. (2009) also propose an extended BDI architecture integrating emotions.

2.3.5 BDI in the scope of this paper

In this section we have explained what is BDI: a paradigm for formalising agent behaviour in terms of mental attitudes (beliefs, desires and intentions). But we have also shown that BDI has multiple implementations, and multiple extensions (with emotions, with norms, etc.). In the end, there are almost as many BDI architectures as there are applications, most of them relying on an *ad hoc* implementation. In the sequel of this paper, we are interested in the methodological and conceptual aspects of the BDI paradigm, by opposition to the engineering aspects. We will call 'BDI' all simulations claiming to use this paradigm, and describing the behaviours of their agents in terms of mental attitudes, whether they use only the original beliefs, desires, intentions or also additional mental attitudes, and whatever their underlying implementation.

3 Methodological guide: when to use BDI agents in simulations?

In this section, we aim at providing a methodological guide to help modellers to choose whether to use a BDI architecture or not, depending on the desired agents characteristics (Section 3.1), on their field of application (Section 3.2), on the goal of their simulation (Section 3.3) or on the target observation level (Section 3.4). We also discuss some arguments against the use of BDI agents, along with possible answers (Section 3.5).

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

3.1 *Influence of the characteristics of the agents*

3.1.1 *Representation and reasoning*

The BDI framework is a model of how humans represent their perceived environment in terms of mental attitudes, and how they reason about it to deduce new mental attitudes from their existing ones. Works in modal logic have investigated these concepts in details, and provided formalisations of their links with each other and of various deduction rules. Modal logic can be seen as a strict and constraining formalism, but it has the advantage of being totally unambiguous.

The BDI model allows a subjective representation of the environment in terms of beliefs, that can be incomplete, erroneous (due to flawed perceptions), or differ between agents. For example, in the AMEL model of evacuation after an earthquake, each agent has their own map of the roads and position of obstacles (Van Truong *et al*., 2013). To account for the uncertainty of perceptions in unknown or partially known environments, Farias *et al*. (2010) have extended BDI agents with fuzzy perception.

Another major benefit is the possibility for the agent to represent and to reason about other agents' mental attitudes (Theory of Mind, Goldman, 2012). For example, Bazzan *et al*. (1999) have developed a model of traffic in which (social BDI) drivers can reason about other drivers' behaviour when making their decision. Similarly, Bosse *et al*. (2007) have developed a BDI model allowing agents to reason about other agents and their reasoning process, and use it to represent a manager's manipulation behaviour over his employees.

Such a capacity to represent and reason about the world in terms of mental attitudes can be used to develop agents that can explain their behaviour in understandable terms to human modellers and users: these are known as self-explaining (Harbers, 2011; Fähndrich *et al*., 2013).

3.1.2 *Self-explanation of behaviour*

The issue of explaining a reasoning is as old as expert systems (Swartout *et al*., 1991), a field where a lot of work has been done to develop explanatory systems and methodologies. In particular, when an expert system is used in a learning environment (for instance, a military training environment (Gomboc *et al*., 2005)), it is very important for the users to be able to get explanations about the reasoning made by the system, and about how and why it behaved as it did. More generally, Ye and Johnson (1995) have shown that an expert system explaining its response was better accepted by end-users: indeed if the system can justify its choices, then it will increase the end-user's trust. This issue is still current in MAS (Fähndrich *et al*., 2013).

In agent-based simulations, the problem is quite similar: a simulator used to teach something is more effective if it is able to give explanations of its behaviour. More generally, Helman and Bahuguna (1986) have shown that explanations of computer science simulations are useful for modellers to better understand the simulation. For example, we can assume that an agent-based simulation used as a decision-support tool by policy-makers (e.g. in urban planning, ecological risk assessment, etc.) would be better understood and accepted, and as a consequence more efficient, if some agents (and in particular those representing human beings or institutions) could be asked for explanations.

The introduction of BDI agents in the simulation is particularly interesting from this point of view. For example, Harbers *et al*. (2010) have investigated the usefulness of BDI agents to provide explanations to the learner in a training simulation. Indeed an interesting capability of BDI agents is that they can explain their behaviour in terms of their goals (why they did what they did) and their beliefs (what were the reasons for their choice of action); and these explanations can become much more precise by adding complementary concepts (e.g. norms) to the basic BDI model. Moreover, the concepts involved in these self-explanations are very intuitive for human beings because they match their own natural language and their way of thinking (Norling, 2009), making explanations provided by BDI agents more effective.

Therefore, BDI concepts facilitate the creation of explanations. Nevertheless, the counterpart of this gain in terms of explanation power is a much more complex reasoning process. This is not necessarily a drawback against the use of BDI agents, on condition that these agents are systematically endowed with a self-explaining capability. Indeed, as highlighted in Harbers *et al*. (2010), without providing this capability to agents, the agents behaviour and therefore the whole simulation become almost impossible

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

to understand, both by computer scientists implementing it and by modellers trying to use it. To put it differently, a side-effect of the introduction of BDI agents in simulations is the necessity to systematically endow agents with a clear self-explanation capability in order to make the simulation understandable, but the implementation of this self-explanation capability is facilitated by the BDI concepts underlying the reasoning of these agents.

3.1.3 Complex interactions via communication

In the agent-oriented programming community, in parallel to the development of the BDI architecture many efforts have been dedicated to the development and formalisation of Agent Communication Languages (ACL) (e.g. KQML (Finin *et al*., 1994) or FIPA-ACL (FIPA, 2002a) to cite the most important ones) and interaction protocols (IPs) (e.g. the FIPA Contract Net Interaction Protocol (FIPA, 2002b)). These ACL and protocols enable structured interaction between agents[2].

Such ACLs and IPs have two main advantages, that can be directly applied to simulation. First of all they allow agents to communicate in a realistic (human-like) way, with a large choice of speech primitives. The communication is closely linked to the mental states of the agents by the ACL semantics (e.g. the FIPA-ACL one (FIPA, 2002a)). Therefore, this allows one to easily model the effect of the communication on the agent behaviour: received messages modify the mental states of the agent, which in turn have an influence on its behaviour. For example, Fernandes & Nunes (2008) propose a traffic simulator in the context of Intelligent Transportation Systems, with communication between vehicles. Thus, the simulator must take into account this communication and its influence on the vehicle driver's behaviour. The authors follow the Prometheus modelling methodology (Padgham & Winikoff, 2002) to integrate BDI agents and communication via FIPA-ACL.

The second interest of using complex communication in the simulation is the use of the IPs. Some protocols have been formalised, in particular by the FIPA consortium, in order for agents to coordinate their tasks. This is particularly useful when the simulation manages a swarm of agents such as in rescue simulations. Some simulation platforms integrate these mechanisms for this purpose. For example Sakellariou *et al*. (2008) have developed a forest fire simulator where firemen agents cooperate and coordinate under the Contract Net Protocol. To this purpose, they have extended the NetLogo platform (Wilensky, 1999) with a tool for easily integrating BDI agents exchanging FIPA-ACL messages in a FIPA protocol. In addition, Thabet *et al*. (2009) show on a grid scheduling example that a MAS involving BDI agents communicating with high-level FIPA-ACL can bring benefits on two levels: at the micro-level, agents are able to perceive the network, discover failures and reason about them to choose the most adapted behaviour; at the macro-level, agents use high-level IPs, providing them with a powerful and flexible way to communicate.

Complex communication mechanism with BDI agents are also used to model negotiation behaviours. Such models have been used, for instance, to simulate negotiation in mix human-agents teams in virtual worlds for training (Traum *et al*., 2003). The aim of this simulation is to improve the leadership skills of human beings by interacting with virtual teammates as believably as possible.

To conclude, when using BDI agents, the use of complex communication with ACL and IP is quite easy, in particular because both share the same philosophical ground and the same concepts. This provides, when it is needed in the model, powerful coordination, cooperation or negotiation mechanisms via communication.

3.1.4 Planning

Plans are sequences of actions that an agent can perform to achieve its goal(s). In some simulation scenarios a modeller may want agents to continuously rethink their behaviour, based on their situation and goal(s), in a manner that cannot be captured by simple if-then rules. This may involve, for example, modelling a complex decision process, supplying the agent with a comprehensive library of plans, or equipping an agent with fundamental planning capabilities.

[2] Agents can also communicate directly via the environment; low-level problems such as communication noise are out of the scope of this paper.

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

Examining the application of multi-agent simulation even in just one domain (crisis response) we already find agents equipped with a broad range of planning capabilities. PLAN C[3] uses reactive agents with simple scripted plans to model medical and responder units in a mass casualty event simulations to explore the medical consequences of possible response strategies. DrillSim (Balasubramanian *et al*., 2006) model fire warden and evacuee agents with a recurrent artificial neural network to select their course of action and an A* algorithm for low-level path planning; these agents are used to test the benefits of information collection and dissemination in a fire evacuation scenario. D-AESOP (Buford *et al*., 2006) models medical relief operations in crisis situations, in an attempt to support more effective organisation, direction and utilisation of disaster response resources. It appears to be necessary to consider the influence of significantly more information (social, medical, geographical, psychological, political and technological) on agents' behaviour. Therefore, the developers claim to need a rich agent model and choose the BDI paradigm where they replace event/goal plans with 'situation' plans that capture a broader view of the disaster situation.

As discussed in Section 2.3, the standard BDI paradigm uses a predefined library of plans tagged by their applicability context, and the goal(s) which they seek to achieve. There exist different extensions to the BDI paradigm for equipping agents with fundamental planning capabilities, so they can build their own plans during the simulation: classical planning (e.g. de Silva *et al*., 2009) lets agents synthesise courses of action depending on available domain knowledge; Hierarchical Planning (e.g. Sardina *et al*., 2006) supports the generation of high-level plans based on hypothetical reasoning about the relative value of plans.

There are some examples of such extensions being used in the context of multi-agent simulation, Shendarkar *et al*. (2008) use agents with a real-time planning capability in a crowd simulation. Schattenberg present JAMES (Schattenberg & Uhrmacher, 2001), a Java-based agent modelling environment for simulation which supports the development of BDI agents with fundamental planning capabilities. The authors claim that 'adding planning capabilities to an agent architecture offers the advantage of an explicit knowledge representation combined with a stable and robust decision procedure for the strategic behaviour of the agent'. Novák *et al*. (2013) work on autonomous military robots teams for surveillance missions: their first implementation using a reactive hierarchical planner was shown to be inflexible due to its extreme simplicity, and thus later replaced with BDI agents. However, these authors noted that BDI agents were too heavy-weight and should be limited to scenarios where complex deliberation was actually needed.

3.1.5 Emotional abilities

Cognitive psychological theories (Ortony *et al*., 1988; Lazarus, 1991) describe emotions in terms of mental attitudes; for example, joy is the conjunction of a desire and a belief that it is satisfied. This straightforward connection with the BDI model played in favour of the use of such emotion theories in agents, as noted by Nair *et al*. 'one of the appeals of cognitive appraisal models as the basis of computational model is the ease with which appraisal can be tied to the BDI framework often used in agent systems' (2005). Conversely, this easy link between BDI and emotional concepts also constitutes an argument in favour of using BDI agents in simulations where emotional abilities are needed.

The standard BDI framework needs to be extended to integrate emotions (see Section 2.3.4), but it then helps to reason about them and understand their influence on the agents behaviour. Therefore, a number of BDI models of emotions have been developed (Adam, 2007; Steunebrink *et al*., 2008; Adam *et al*., 2009), and a number of BDI agents have been endowed with emotions to make them more believable (e.g. embodied conversational agents (de Rosis & *et al*., 2003; Guiraud *et al*., 2011)) or efficient (e.g. virtual tutors (Gratch & Marsella, 2005)). Such emotional agents also find applications in various types of simulations, for example, to model panic in an evacuation scenario (Minh *et al*., 2012) or to model the population possibly irrational behaviour in a bushfire (Adam *et al*., 2015).

In addition to the appraisal process (the way stimuli are appraised to trigger an emotion), the coping process (the way agents deal with their emotions via various strategies) is also very important as it impacts

[3] `http://www.nyu.edu/ccpr/laser/plancinfo.html`.

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

the agents reasoning and behaviour, and as such has already been integrated in some BDI agents (Adam & Longin, 2007, 2014). For example, Gratch and Marsella (2004) have developed a general BDI planning framework integrating appraisal and coping processes to trigger emotions in an agent, based on its interpretation of the world, and to model their influence on its behaviour. This framework called emotion and adaptation (EMA) has been applied to military training simulations: human officers interact in a virtual world with virtual humans in order to improve their skills in real-life stressful situations. Complex agent behaviours (in particular emotional behaviours) are needed to add realism and favour immersion, which in turn enhances training.

BDI agents have also been endowed with emotions in simulations aiming at exploring the relationships of emotions with various other concepts. Staller and Petta (2001) show that emotions have a bi-directional relationship with norms: emotions can be an incentive to respect norms (e.g. shame or guilt) or to enforce them (e.g. contempt); and norms influence the triggering and regulation of emotions. They have proposed an implementation in TABASCO and aim at reproducing the results of Conte and Castelfranchi (1995). Nair *et al*. (2005) have investigated how emotional abilities can improve teamwork in teams composed of either only agents or of both agents and humans; agents simulating human beings have a BDI architecture. Inspired by neurosciences, Bosse *et al*. (2013) modelling proposes an evacuation simulation where the agent behaviour model is based on beliefs, intentions and emotions, and the relationships between these three concepts. Authors also take into account social interactions, and in particular emotional contagion. Jones *et al*. (2009) have developed a BDI extension called PEP → BDI that enriches the BDI engine with personality, emotions, as well as physiological factors; they apply it to a simulation aimed at training security actors to handle major scale crisis, such as terrorist attacks.

However, not all simulations use BDI agents to model emotions, some use very simple mechanisms instead to deal with emotions. For example, in Minh *et al*. (2012), emotion is represented as a single numerical value that increases with perceived danger and influences from others, and decreases with time. Authors claim it is enough to deal with the impact of emotions on evacuation in case of a crisis. But much more complex mechanisms can also be used. For instance, Silverman *et al*. (2005) have integrated emotions in their multi-agent communication for evacuation simulation (MACES) crowd simulation system by using PFMserv, a cognitive architecture claimed to enable more realistic agent behaviour, by providing the agents with a working memory similar to human short-term memory, as well as a long-term memory, both used in their decision cycle.

Although other (more simple or more complex) models of emotions exist, the BDI model brings major benefits in social simulations, because it allows reasoning and interpreting emotions of self and others (reverse appraisal, Adam *et al*., 2011; Rivière *et al*., 2012), and formalising their influence on behaviour via coping strategies.

3.1.6 Reasoning about norms

Norms are a set of deontic statements (obligation, permission or prohibition) specifying how agents are expected to behave; they are often represented in deontic logic (von Wright, 1963; Conte & Castelfranchi, 1995). Social norms are norms shared by a group of agents; they are a very powerful tool to describe the rules in a group of agents. It is important to consider the fact that these norms *can be* violated; therefore a mechanism to deal with violations is needed, such as sanctions applied to guilty agents. Each agent has to take norms into account in his decision-making process, his desires should be balanced by norms. The basic version of BDI does not integrate norms, but some extensions do, for example, the BOID architecture (Broersen *et al*., 2001; Neumann, 2010) which focusses particularly on conflicts between the four modalities.

Norms are very important in several applications of MAS (such as business-to-business or e-business) but they are particularly essential in the modelling of socio-environmental systems (Ostrom, 2007; Gaudou *et al*., 2013). Savarimuthu and Cranefield (2009) draw a quite complete overview of the use of norms in simulation, and categorise simulations *w.r.t.* their normative mechanisms. In simulations where the norms and their influence on the agents behaviour are of interest, it is important to have a fine-grained model of agents' cognition that allows agents to explicitly reason about norms (see Section 3.4). For example, some simulations aim to cognitively explain criminal behaviour (Bosse *et al*., 2009) and thus use BDI agents.

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

In Conte and Castelfranchi (1995), the function of social norms is to control aggression in society. The authors deliberately used very simple agents to study the function of norms at a macro-level. These agents can perform some simple routines (move, eat, attack) to survive in an environment with few food resources; they have a strength that is increased by eating and decreased by movement and attack. The authors compared two strategies to control aggressions in this artificial society: in the normative strategy the agents respect a 'finder-keeper' norm, while in the utilitarian strategy they attack for food any agent not stronger than themselves. The simulation showed that the normative strategy was much better to reduce the number of aggressions. Staller and Petta (2001) extend Conte and Castelfranchi (1995) but focus on studying the bi-directional relationship between emotions and social norms: social norms influence the triggering of emotions (e.g. feeling ashamed when violating a norm) and their regulation and expression (e.g. one should not laugh at a funeral); and emotions in turn influence normative behaviour (e.g. shame drives people to respect norms, while contempt drives them to apply sanctions to enforce norms). Due to their focus on the micro-level and its interactions with the macro-level, the authors cannot content themselves with the very simple agent model used in Conte and Castelfranchi (1995). Instead they use the TABASCO three-layer agent architecture, and a BDI architecture called JAM for the reasoning part of their agents.

BDI architectures seem very adapted to model two mechanisms: the internalisation of norms, and the influence of norms on the decision-making process. Indeed, the explicit representation of norms allows an agent to reason about them and make autonomous decisions regarding their internalisation (instead of automatically accepting all norms); based on its desires and beliefs, the agent can thus choose to violate norms (such as in Castelfranchi *et al*. (2000).

One of the main critics against BDI architecture regarding norms is that such architectures use *a priori* existing norms (implemented offline). The emergence and recognition of norms at runtime can be handled with other architectures such as EMIL (Campenn *et al*., 2009); but this problem has only been marginally and very recently tackled in the BDI literature. Criado *et al*. (2010) extend a graded BDI architecture to allow autonomous (and agent-dependent, based on their preferences, contrarily to EMIL) norm acceptance.

To conclude, BDI architectures offers a powerful tool to reason on norms and to integrate them in the decision process, but are more adapted to simulations with initially established norms. Therefore, the emergence of informal social norms in groups of agents would be harder to tackle with a BDI architecture. However, this restriction is not so strong since BDI architectures are still applicable to the large number of simulations where laws and regulations are established by a superior institution.

3.1.7 Decision-making

First of all, we can notice that the decision-making process is often not key in BDI frameworks and architectures. Norling (2009) notices that BDI theory assumes that agents make utility-based rational decisions, optimising the expected utility of the selected action, and that in implementations (e.g. JACK) agents usually choose the first applicable plan regardless of its expected utility. But Norling claims that both strategies are not representative of human decision-making, and proposes to extend BDI with a different decision strategy: recognition-primed decision (RPD). This strategy is typically used by experts who spend most of the time finely analyzing the situation, before making an almost automatic decision based on this situation.

Some works show that the BDI paradigms eases the modelling of agents with a more realistic decision-making process. Lui *et al*. (2002) model military commanders in land operations scenarios with a BDI approach, which they claim to provide the agents with the ability to demonstrate human-like decision-making. Noroozian *et al*. (2014) use BDI agents to model individual drivers with heterogeneous driving styles for a more realistic traffic model. Indeed, the BDI paradigm makes it easier to design agents with various decision criteria that can react differently to the same stimulus.

In addition, the BDI framework makes it easy to integrate and explore the influence of several cognitive factors on the decision-making process. Bosse *et al*. (2007b, 2009) present a BDI modelling approach to criminal decision-making in a violent psychopath agent. Their model is based on criminology theories that

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

define opportunities for a crime. Desires are generated from biological, cognitive and emotional factors; beliefs in opportunities are generated from the observation of the social environment; their combination leads to intentions and actions to satisfy these desires. The implementation is done with the authors' LEADSTO language that allows them to combine quantitative (physiological hormones levels) and qualitative (mental attitudes) concepts in a unified model. In a socio-environmental application, Taillandier *et al*. (2012) build a model to deal with farmers' decision-making about cropping plans, farmers behaviours being described with a BDI architecture. To this purpose, they have used a multi-criteria decision-making process, considering the desires as the criteria that drive the decision. Finally, Swarup *et al*. (2014) argue that a BDI model of agents can be fully relevant to understand and model seemingly irrational human decision-making, such as H1N1 vaccination refusal, and its influence on the dynamics of an epidemic.

3.1.8 Summary

To summarise, the BDI model is quite agnostic in terms of the decision-making approach, and many algorithms can be used. Nevertheless, modellers can benefit from the high-level BDI concepts to design a much finer decision-making process than the simple economic maximiser agent. The BDI framework also allows to model individualised, possibly irrational, decision-making, where each agent makes its decision based on its own motivations. This is essential to simulate realistic human societies for various goals in various fields of application and for various goals. Table 1 summarises the importance and benefits of using a BDI architecture for each of the characteristics above.

3.2 Influence of the field of application

Macal and North (2005) in their tutorial on ABMS list a number of fields in which there are practical applications of ABMS. We examine each of these fields as well as some other more recent ones, and discuss the appropriateness and benefits of a BDI approach in each case.

In business and organisations simulations: most *manufacturing* applications involve clear rules and lists of tasks, so a goal-based approach like BDI is not the most appropriate; on the contrary, BDI offers modellers of *consumer markets* a straightforward mapping between their understanding of the purchasing decision-making process and the BDI concepts (see e.g. Baptista *et al*., 2014); finally in *supply chains and insurance* simulations, depending on the complexity and uncertainty in the models, BDI may also help as it allows a more complex but still transparent model.

In economics, *artificial financial markets and trade networks models* are presently often over simplified by assuming perfectly rational agents. By using BDI agents, emotions and norms, a more reasonable assumption of the human reasoning process could be captured (Wan *et al*., 2013).

In *infrastructure* simulations, the participants of *electric power markets or hydrogen economies* need to 'perform diverse tasks using specialized decision rules' (Macal & North, 2005), which could be achieved by specifying a plan library and allowing the BDI paradigm to capture the complexity. Nowadays with power grids, the electric power market model needs to integrate micro-behaviours of inhabitants to be more accurate; BDI agents are relevant to simulate 'social, perception, cognitive and psychological elements to generate inhabitants' behaviour' (Kashif *et al*., 2011), based on a reasoning about potential consequences of their actions to find the most energy-efficient behaviour.

In *transportation simulations*, BDI is ideal for modelling humans who have many different ways of using the transportation systems, each with its own benefits. The self-explanatory nature of BDI agents also allows the modeller to understand why one choice was made over another. It can also help to model individual drivers with heterogeneous driving styles (Noroozian *et al*., 2014) for a more realistic traffic model.

In crowds simulations, in particular *human movement and evacuation modelling*, BDI offers more precise and realistic modelling of human behaviour, in order to enrich the crowd simulation (Cho *et al*., 2008; Shendarkar *et al*., 2008). For instance, Park (2013) made a crowd simulation more realistic by introducing social interactions between the agents.

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

Table 1 Summary of benefits and drawbacks of BDI architecture agents characteristics (learning will be discussed in Section 3.5.4)

Agent characteristics	Importance of BDI	Benefits and drawbacks
Representation and reasoning	+++	Non-ambiguous representation
		Theory Of Mind (reasoning about others)
		Limited basic expressivity overcome by extensions
		Complex reasoning: small-scale simulation
Self-explanation	+++	Intuitive explanation (folk psychology)
Communication	++	High-level communication (coordination)
		Mentalistic semantics of ACL
		Flexible interaction protocols
Planning	++	Stable, robust decision procedure
		Complex (heavy) but flexible deliberation
Emotions	+	Formal description, links with mental attitudes
		Reasoning and interpreting (reverse appraisal)
		Influence on behaviour (coping strategies)
Norms	+	Explicit representation allows reasoning
		Integration in decisions-making
		Handling of pre-established norms only
Decision-making	−	Heterogeneous decisions, inter-agent differences
		Irrational decisions, influence of psychological factors
		Practical reasoning
Learning	−	Symbolic learning, not standard in BDI

ACL = Agent Communication Languages.

More generally, *crisis management* needs tools for decision-support for the stakeholders, training of the risk managers, and education and awareness of the population. Such tools require realistic models for better immersion and for valid results. For example, Paquet *et al*. (2004) model rescue workers after a natural crisis in DamasRescue using BDI agents; Buford *et al*. (2006) develop a Situation-Aware BDI Agent System for Disaster Situation Management.

In *epidemiology*, Swarup *et al*. (2014) present the BDI model as very good for understanding seemingly irrational human decision-making such as H1N1 vaccination refusal, and their influence on the dynamics of an epidemic.

In society and culture when simulating the fall of *ancient civilisations* such as in Axtell *et al*. (2002) and Heckbert (2013), BDI agents are irrelevant for several reasons: the time scale of the simulation (several decades); the lack of precise behaviour data; and the focus on external influences (environment, trading) rather than individual decision-making. Nevertheless, BDI agents could be useful when simulating shorter and more recent historical periods where more data is available (e.g. the 1926 flood management in Hanoi (Gasmi *et al*., 2015)). The concepts of norms, as well as the heterogeneity of behaviours among people in an organisation, also make BDI appropriate for modelling *organisational networks*.

In military simulations, since BDI reflects the way people think they think (meta-reasoning), it is perfect for simulating people in a *participatory command and control simulation* (Lui *et al*., 2002; Gratch & Marsella, 2004). The BDI approach has also been used in this field for the realistic decision-making it brings. Karim and Heinze (2005) propose an autonomous drone pilot with BDI reasoning; Novák *et al*. (2013) use BDI reasoning for autonomous surveillance robots in urban warfare, because it is more flexible than planning and more practical than tele-operation.

In *criminology*, BDI agents are used as they can manage various cognitive concepts in a unified architecture. Bosse *et al*. (2007b, 2009) investigate the influence of psychological (emotions) as well as social and biological aspects on criminal behaviour; their model merges qualitative BDI aspects and quantitative biological aspects.

In *ecology*, the lack of a decision-making process in the agents being modelled suggests that BDI is not appropriate. In biology, for *animal group behaviour*, BDI may be appropriate, depending on the animal

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

being modelled and its cognitive abilities; for *cells or sub-cellular molecules simulations*, the simple nature of the behaviours makes BDI unsuitable.

Intuitively, BDI agents are more adapted in fields where the agents modelled are humans and the decision process is well known, and less adapted to the modelling of simpler non-human entities or when there is no decision-making process.

3.3 *Influence of the goal of the simulation*

ABMS consists of two steps: first, create an agent-based model of a target theory or from observed data; second, build and run a simulation of this model. Axelrod (2005) admits that depending on the goal of an agent-based model one should emphasise different features. Concretely, when the goal is to enrich our understanding of a phenomenon, the model has to stay simple to facilitate analysis of the results: 'when a surprising result occurs, it is very helpful to be confident that one can understand everything that went into the model'; it also allows easier replication or extension of a model. On the other hand, when the goal is to faithfully reproduce a phenomenon, and realism and accuracy are important, the model might need to be more complicated: 'there are other uses of computer simulation in which the faithful reproduction of a particular setting is important […]. For this purpose, the assumptions that go into the model may need to be quite complicated'. Once one has designed a model, appropriately balancing accuracy and simplicity, the next step is to run a simulation, which can be used to reach various goals. Indeed Axelrod (2005) lists seven purposes of simulations. In the following paragraphs we present each of them, provide examples of simulations having this goal, and discuss if BDI agents are appropriate to reach it or not, and why.

3.3.1 *Prediction*

Simulations can be used to predict potential outputs resulting from known complex inputs and assumptions about mechanisms.

For example, Gunderson and Brown (2000) use data mining to create a simulation from a large quantity of statistical data available about crimes, and then use it to predict criminal behaviour. The authors are not interested here in the details of the agents cognition, but in the global crime pattern at the scale of a city, and therefore do not need BDI agents (nor have the micro-level data needed to design a BDI model, plus it would add unneeded complexity and scalability problems). In a different field, Baptista *et al*. (2014) model consumers with a realistic underlying BDI agent model, which they claim to be valuable for predicting market trends.

Simulations that can predict the behaviour of the population in response to various policies are also a powerful tool for decision-support, on condition that the underlying model is as faithful as possible in order to lay valid results. A BDI architecture of agents representing humans thus ensures the right level of realism. Indeed, Swarup *et al*. (2014) in their prospective paper claim that the BDI model is very well adapted to model (possibly irrational) human decision-making, such as H1N1 vaccination refusal, when simulating and predicting the dynamics of an epidemic.

Such decision-support systems can be used to assist policy-makers in various fields. For example, Sokolova and Fernández-Caballero (2007) present a decision-support system that analyses pollution indicators and evaluates the impact of exposure on the health of the population, modelled as BDI agents; their system allows ecological policy-makers to explore what-if scenarios, giving them evidence to support decision-making. Shendarkar *et al*. (2008) propose a crowd model with BDI agents to simulate evacuation in case of a terrorist attack, in order to determine where it is more useful to place policemen to improve evacuation.

To conclude, the BDI architecture is a good candidate in simulations aimed at predicting future behaviour, because this requires the behaviour of agents to be as realistic as possible, adaptable to a large number of situations. However, it can only be used if there is enough detailed data available: this will be discussed in Section 3.5.1.

3.3.2 *Performance of a task*

Axelrod (2005) claims that 'to the extent that artificial intelligence techniques mimic the way humans deal with these tasks, the artificial intelligence method can be thought of as simulation of human perception,

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

decision making, or social interaction'. We therefore sort in this category artificial intelligence research that attempts to simulate human processes into agents to improve their efficiency at these tasks.

For example, the Virtual Humans community aims at creating artificial humans endowed with a wide range of human-like abilities (including face-to-face multimodal interaction, emotions, personality, etc.) (Gratch *et al*., 2002; Swartout *et al*., 2006). Agents endowed with good (human-like) communication skills are useful as interface agents or as virtual companions (Wilks, 2010).

When virtual humans aim at a more realistic or complex behaviour, they usually rely on a BDI architecture, possibly extended with additional required features. For example, BDI agents with human-like negotiation skills have been used as mentors in role-playing games to train humans to improve their own negotiation skills in stressful situations (Traum *et al*., 2003). Another BDI agent endowed with emotions could express realistic emotions or to detect the user's ones in Ambient Intelligence scenarios (Adam *et al*., 2011). Emotions and associated coping strategies were also used as human-inspired heuristics in a BDI reasoner for efficient robot (Adam & Lorini, 2014).

BDI agents with human-like emotional abilities have also been shown to perform better teamwork. Nair *et al*. (2005) use a BDI model of coordination, combined with POMDP to handle uncertainty. In pure agents teams, their BDI emotional agents have a more realistic behaviour, since humans are known to use emotions for goal prioritisation and action selection. In mixed humans-agents teams, they are better teammates: awareness of their human partners' emotions improves communication with them, and expression of their emotions contributes to building a relationship with them.

Peinado *et al*. (2008) present BDI as a 'promising tool for modelling sophisticated characters in Interactive Storytelling'. They mention in particular that it allows to create richer characters with their own beliefs and motivations causing their behaviour

Karim and Heinze (2005) use BDI agents to design an autonomous pilot for military unmanned aerial vehicles (UAV); they find that the BDI paradigm provides a more modular design, a high level of abstraction, and improves scalability by avoiding hard-coded coordination.

So BDI agents seem to provide a good paradigm to realistically model human behaviour in agents aimed either at behaving in a believable, human-like way, or at performing tasks more efficiently by reproducing the human way of performing them. Indeed BDI provides a high-level description of such behaviour, so it is the perfect level of abstraction to efficiently capture psychological theories of human behaviours. Such agents endowed with human-like abilities find a very wide variety of applications, from virtual companions or storytellers to military UAV pilots.

3.3.3 Training

These simulations aim to 'train people by providing a reasonably accurate and dynamic interactive representation of a given environment' (Axelrod, 2005). Axelrod then notices that simulations having such a goal need accuracy and realism rather than simplicity. Indeed lots of training simulations do use complex BDI agents, often extended with additional abilities such as emotions or negotiation skills.

A lot of works have used BDI agents in serious games for training, because they provide a human-like decision-making process described in terms of folk psychology. Baptista *et al*. (2014) thus model consumers with a realistic underlying BDI agent model, and claim it can be useful for training management skills of business students. Another application is the use of virtual humans as virtual patients to train medical professionals. Rizzo *et al*. (2011) have used a virtual human with a complex BDI architecture to create different virtual patients that provide realistic representations of mental problems to be diagnosed or discussed. Such virtual patients are useful to improve doctors' interview and diagnostic skills for a difficult subject through practice.

The same virtual human has also been used in a variety of other simulations for the military. For example, it has been used in a war scenario, the Mission Rehearsal Exercise, to train military task-negotiation (Traum *et al*., 2003) and decision-making (Gratch & Marsella, 2004) skills. Traum then enriched it with conversation strategies (Traum *et al*., 2005) to create a virtual negotiator that can play the role of a mentor to train military negotiation skills in crisis situations; the trainee learns principles of negotiation by interacting with the agent using them to play the opposing party. Many other military

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

applications can also been found. Mcilroy and Heinze (1996) developed the Smart Whole AiR Mission Model that splits the physics of combat and pilot reasoning into two separate entities which are modelled, validated and implemented in completely different ways; fighter pilots have a BDI architecture based on the distributed multi-agent reasoning system (dMARS) (D'Inverno *et al*., 2004) realisation of the Procedural Reasoning System. Lui *et al*. (2002) use Jack Teams, a BDI team modelling framework, to model military commanders in land operations scenarios.

BDI agents have also been used because of their adaptability to a changing environment. Oulhaci *et al*. (2013) designed a serious game for training in the field of crisis management, where non-played characters (NPC) have a complex behaviour, either behaving as expected or intentionally erroneously. These NPC are developed as BDI agents so that they can adapt to events of the world and to social interactions with and actions of other players (in particular the learner). Cho *et al*. (2008) used BDI agents to simulate crowds in virtual worlds, claiming that BDI increases their adaptability to the dynamic environment, and therefore their realism. They present a prototype game with NPC extinguishing a forest fire, implemented with the SourceEngine three-dimensional (3D) game engine.

To conclude, simulations for training require a realistic immersive environment to be efficient. In particular, the behaviour of autonomous agents populating the virtual world (NPC) must be human-like, making BDI agents a very good choice, for their descriptivity of the human-like decision-making process, and their adaptability to a dynamic environment.

3.3.4 Entertainment

This goal is quite similar to training but the virtual worlds can be imaginary ones. There has been work on integrating BDI agents into games in order to provide the players with a richer, more enjoyable experience.

A number of works are dedicated to the integration of autonomous agents (bots) in Unreal Tournament. Small (2008) uses a simple rule-based agent called Steve, that can quickly decide its action (e.g. move to healing source) based on high-level conditions (e.g. weak energy) in a dynamic environment. To be able to compete with human players, Steve mimicks their learning process through an evolutionary algorithm: its strategies are evolved (and improved) over time. However, it was still unable to consistently beat human challengers. Hindriks *et al*. (2010) wanted to find alternatives to such reactive bots. They argue that using BDI agents instead of scripts or finite-state machines (often used in real-time games) might result in more human-like and believable behaviour, thus making the game more realistic. They add that when one gets data from observing actual game players, it is easier to specify it in terms of beliefs, desires and intentions and import it into BDI agents, than to translate it into finite-state machines. They thus tried to interface the GOAL agent programming language with Unreal Tournament 2004. Such an interface raises two main issues: finding the right level of abstraction of reasoning (the BDI agent should not have to deal with low-level details, while keeping enough control), and avoiding cognitive overload to achieve real-time responsiveness (the agent should not receive too many percepts, but should have enough information, even if possibly incomplete and uncertain, to make decisions). One of their motivations was to test agent programming platforms, in order to eventually make them more effectively applicable. They indeed noticed that BDI agent technology is not scalable yet, and they observed bad performance when increasing the number of agents.

Another try at integrating BDI agents into Unreal Tournament is described in Urlings *et al*. (2006). The authors are actually interested in humans-agents teams in the area of unmanned vehicles. They propose a framework to model these and then use the Unreal Tournament game as a real-time simulation environment to demonstrate this framework. Their agents are implemented with a BDI architecture using JACK. BDI agents have also been used in other games: Norling (2003) describes a BDI Quake player; the main character of the Black and White game (Molyneux, 2001) is also a BDI-based agent. In another attempt to make video games more immersive and realistic (and not only graphically so), Palazzo *et al*. (2013) provide the Spyke3D framework, using BDI agents to facilitate the creation of more human-like NPC with plausible logical behaviours.

The works described above all concern video games, but BDI can also be used to model characters in interactive stories and endow them with a richer causal behaviour based on their beliefs and motivations (Peinado *et al*., 2008).

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

So BDI agents are an excellent tool to create believable characters that make games or stories more realistic and immersive. Unfortunately, there are performance issues when the number of agents increases, which curbs their wider use in games so far, when designers usually prefer to improve graphics and rely and simple but fast scripted behaviours. We will discuss solutions to the scalability problem in Section 3.5.2.

3.3.5 Education

According to Axelrod, 'the main use of simulation in education is to allow the users to learn relationships and principles for themselves' (2005). So this goal is similar to training (and entertainment), but the virtual world does not need to be as accurate (Axelrod says they 'need not be rich enough to suggest a complete real or imaginary world'), as long as it lets the student experiment and learn by trial and error. For example, in the *SmartEnergy* project, Kashif *et al*. (2011) propose a serious game aimed at raising awareness about energy saving; the realism is not key here as the main requirement is to give the users an insight of the effects of their actions.

Nevertheless to improve the users' learning, it is often useful to provide them with an artificial tutor inhabiting this virtual world, and able to explain to them why they did wrong or right. Therefore, Harbers *et al.* (2010, 2011) use BDI agents that can provide such explanations in terms of intuitive concepts explicitly represented in the BDI architecture: their goals and beliefs (see Section 3.1.2). For example, Elliott *et al.* (1999) study the integration of emotions in a virtual tutor called Steve that inhabits a 3D world and teaches students to operate a high-pressure air compressor; Steve has also been used in a training simulation with an army peacekeeping scenario to prepare officers to deal with difficult situations (Rickel *et al*., 2001).

To conclude, in education-oriented simulations realism of the virtual world and of its population is not key, so it is not necessary to use a BDI architecture for all agents. Nevertheless, the BDI architecture is really useful for some specific agents that need to be able to explain their behaviour or the user's errors, such as some NPC or an artificial tutor, in order to improve learning.

3.3.6 Proof

Another possible purpose of simulations identified by Axelrod is to provide an existence proof. For instance, Conway's Game of Life (Gardner, 1970) is a game based on a set of simple rules governing the birth and death of cells, which can be used to prove or disprove conjectures such as 'a finite population cannot grow infinitely large'.

The literature provides few examples of agent-based models aimed at existence proofs. One famous example, however, is Schelling's (1971) agent-based segregation model. Very simple rules govern the agents' movement to other neighbourhoods based on their level of happiness with their neighbours. Just a small desire to live near people with similar traits can quickly lead to large clusters of homogeneous neighbourhoods. This model is considered to provide an existence proof that segregation does not require the people involved to be highly intolerant, but only mildly so.

The line between exploration and proof is difficult to draw. Axtell *et al*. (2002) provide a multi-agent model of the exodus of the Kayenta Anasazi (modelled as relatively simple reactive agents) from Long House Valley, as a means to explore its possible causes. Whilst it provides a strong suggestion that the inability of the land to sustain the population was not the only cause of the exodus (by showing that some subset of the population could have been sustained) it falls short of being considered an existence proof that socio-cultural factors had some influence. It is not clear whether this is because this relatively simple model is in fact too complex, or indeed too basic to support such reasoning.

Sun and Naveh (2007) argue that socio-cognitive simulation is essential to prove functionalist hypothesis, that justify social phenomenons by the function they have in society, that is, that explain a cause by its effect. Such hypothesis are usually hard to verify, but social simulation allows to verify or prove them through experimentation. However, the rudimentary cognition assumed in agents used in these simulations does reduce their realism and hence their applicability, and also prevents from proving links between the micro- (cognitive) and macro- (social) levels of the simulation. The authors therefore argue for the use of a specific cognitive architecture called CLARION; we claim that BDI agents would achieve the same purpose (as discussed above in Section 2.2).

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

To conclude, depending on the hypotheses that the simulation aims at proving, the complexity required from the agents involved is not the same. Hypotheses about the survival of a cell population, or the geographical movements of a human population, can be proved with very simple agents; in contrast, more complex hypotheses about the links between individual cognition and social organisation do require more complex agent models, with BDI being one option.

3.3.7 Discovery

The last purpose of simulations identified in Axelrod's classification is knowledge discovery. Axelrod claims that social scientists have successfully 'discover[ed] important relationships and principles from very simple models' and argues that 'the simpler the model, the easier it may be to discover and understand the subtle effects of its hypothesised mechanisms' (2005).

Sun has tried to discover relationships and interactions between the micro-level (cognition of agents) and macro-level (dynamics of society) (Sun, 2004), for example, in the simulation of social groups depending on the food management strategy (Sun & Naveh, 2007). He obviously needed a fine-grained model of agents, and used his CLARION cognitive architecture. Heinze *et al*. (2001) model airborne early warning and control crew with BDI agents in a simulation used for the exploration, evaluation and development of tactics and procedures, in order to advise the Australian Defence Force in the conduct of maritime surveillance and patrol.

ABMS have also been used recently as a tool for public policy analysis, as it provides a means to discover possible sources of *policy resistance*. Policy resistance is the tendency for policy interventions to be defeated by unpredicted effects. For example, making public fuel-efficient transportation more affordable to the general public may reduce carbon emissions around the city, but the money saved due to reduced cost of local road travel might be spent on leisure, leading to an increase in air travel, resulting in no overall reduction in carbon emissions. Such effects are generally not intuited because the system is constantly changing, adaptive and evolving, with tightly coupled actors and nonlinear effects (Sterman, 2006)—factors which are well addressed by agent-based modelling (ABM).

For example, the Electricity Market Complex Adaptive System multi-agent simulation has been used to explore the restructuring and deregulation of the electric power market in Illinois (Cirillo & *et al*., 2006), and provided information on this policy issue 'that would otherwise have not been available using any other modelling approach' (Macal & North, 2005). There are also many examples of agent-based simulations being used to study climate policy (e.g. Moss *et al*., 2001; Posada *et al*., 2008).

Whilst policy development is an emerging field of application for ABMS, the tight link between policy impact and human behaviour strongly suggests a need for realistic, rich models of the latter, as supported by the BDI paradigm. In the context of using agent simulation for exploring economic policy, Farmer and Foley (2009) suggest that 'the major challenge lies in specifying how the agents behave and, in particular, in choosing the rules they use to make decisions'.

3.3.8 Conclusion

To summarise, we showed in this section that the BDI architecture can bring benefits to simulations whatever their goal, but more particularly so for some goals. Simulations for *prediction* needs a more realistic underlying model for all agents (with BDI being a candidate of choice for human agents) to accurately support decision and to lay valid results. When simulations aim at imitating human processes for *performing a task*, BDI provides the right abstraction level to capture the modelled reasoning. Simulations aiming at *training, entertainment or education* benefit from realistic and adaptive agents, that enhance immersion and engagement, and increase efficiency by providing explanations; BDI agents are already widely used in such simulations. Finally, simulations for *proof and discovery* should remain simple at first, but BDI agents can then be useful to prove social phenomena involving individual cognition of agents. Table 2 summarises the interest, benefits and drawbacks of using the BDI approach depending on the simulation purpose.

Independently of the goal of the simulation, the choice of a BDI architecture for the agents also depends on other factors, such as the required level of granularity or observation scale (does the simulation focus on

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

Table 2 Summary of interests of BDI *w.r.t.* agents characteristics

Simulation purpose	Interest of BDI	Benefits	Drawbacks
Prediction	Medium	Realistic, adaptable micro-level behaviour possibly irrational individual cognition	Complexity, scalability Requires detailed data
Performance of a task	High	Right level of abstraction of human-like behaviour Awareness, cooperation in mixed human-agent teams Modular, scalable, flexible design	Harder design, uncommon paradigm
Training	High	Accurate realistic behaviour for better immersion Adaptability to a dynamic environment Descriptivity	Harder design, uncommon paradigm
Entertainment	Medium	Rich plausible human-like behaviour: immersion, challenge Quick decision under uncertain and incomplete information Right abstraction level to capture real players' strategies	Scalability, performance Harder design vs. scripts
Education	Medium	Intuitive explanations of behaviour via built-in folk psychology concepts (BDI)	Unneeded realism and complexity for non-essential agents
Proof	Low	Realistic cognition required to prove micro–macro links and complex socio-cognitive phenomena	Unneeded realism and complexity to prove simpler hypothesis
Discovery	Low	Realistic fine-grained behaviour model to discover non-intuitive effects and micro–macro links in adaptive dynamic complex systems	Harder understanding and deduction Harder specification of decision rules

the micro- or macro-level?), and the (type and quantity of) data available to design the model. These will be discussed below.

3.4 Influence of observation level

In this section we argue that the desired level of observation, that is, the focus on micro- vs. macro-level phenomena, influences the choice of the agents architecture. Models in which the agent is the centre of interest seem to be more likely to use a BDI architecture, since the modeller tends to endow the agents with as realistic a behaviour as possible (for instance, for NPC in video games (Norling, 2003), as discussed in Section 3.3.4). This is particularly true in models designed to investigate in details the effects of psychological or physiological factors on individual cognition (for instance, in criminal behaviour (Bosse *et al*., 2007,a,b) as discussed in Section 3.2).

On the contrary, simulations that are focussed on the dynamics of a large population (e.g. crowds simulations or population dynamics) do not need such a fine granularity, and thus generally use simple agents to reduce computing time and enhance performances (Wolfe *et al*., 2008). For example, a lot of crowd simulations use distributed behaviours producing a global flocking behaviour (Reynolds, 1987). However, as soon as modellers want to simulate a more realistic crowd, for example, to improve immersion of the human player in a virtual environment, they try to use a BDI architecture (Silverman *et al*., 2005; Cho *et al*., 2008; Shendarkar *et al*., 2008). The focus of observation has shifted from a global point of view on the overall behaviour of the crowd (e.g. flocking) to a local point of view on the behaviour of individual agents composing the crowd.

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

For example, in criminology, some authors focus on global crime patterns at the scale of a city and use statistical data in a simple agent behaviour model (Gunderson & Brown, 2000), while others investigate the effects of psychology and physiology on individual criminal behaviour with a more complex agent model merging qualitative BDI and quantitative biological aspects (Bosse *et al*., 2007a, 2009). We could also cite traffic simulation, where very simple models of the traffic flow exist (Nagel & Schreckenberg, 1992), but some authors recently used a BDI model of drivers to focus on their individual driving styles (Noroozian *et al*., 2014).

We can conclude that the level of observation of the simulation has a great influence on the choice of an agent architecture: if the focus is on observing the global behaviour of the group of agents, then simple agents are sufficient and should be used to optimise performances of the simulator; but when the individual agents in the group become the focus, a more fine-grained and realistic approach has to be favoured, and the BDI model is more appropriate. However, such an approach also requires more fine-grained data to elaborate the model; this is discussed below in Section 3.5.1.

3.5 Features of a simulation that deter from using BDI agents

In this section, we discuss the factors traditionally cited as limits of the BDI architecture and hindering its wider use in social simulations, along with possible solutions to overcome these limits. We discuss the lack of fine-grained data to design the model (Section 3.5.1), the scalability issue (Section 3.5.2), the limited expressivity (Section 3.5.3) and the absence of learning (Section 3.5.4).

3.5.1 Influence of the availability of data

To build an model as realistic and useful as possible, the data available about the reference system has a huge importance, and can even influence the choice of the agent architecture. Indeed, if the only data available describes the behaviour of social groups (e.g. family statistics) the modeller does not have enough precise information to give a very rich architecture to the agents representing the individual members of these groups (Gil-Quijano *et al*., 2007). For instance, Gunderson and Brown (2000) use data mining to deal with the large quantity of statistical data available about crimes; this lets them model the global crime pattern at the scale of a city, but not the details of the agents cognition. Therefore, when the goal of the simulation (as discussed above in Section 3.3) requires realism, in particular when the simulation is used for prediction, a great deal of attention has to be paid to the collection and processing of precise data, and thus the use of efficient and dedicated techniques and tools.

It is generally much easier to collect objective and measurable data (such as age, income, etc.), which is sufficient to build a simple reactive agent model, than to gather subjective information explaining behaviour in terms of (not observable) mental states. But since this kind of subjective information is needed to build rich agent models such as BDI agents, another approach is to use participatory modelling: put a human in a virtual world to capture his behaviour, and then reproduce it in the agents of the simulation (Shendarkar *et al*., 2006, 2008). For instance, Silverman collected attributes values governing BDI characteristics by placing human beings into a Virtual Reality environment, to build a BDI model of the crowd behaviour after a terrorist attack (Shendarkar *et al*., 2008). The MAELIA project (Gaudou *et al*., 2013) Maelia investigates policies for water resource management; as the water used in the studied basin is mainly influenced by farming practices, the modellers interviewed farmers about how they made their decisions regarding cropping.

So the BDI model needs precise data about individual behaviours, which is harder to collect. But the BDI model also offers advantages, in particular its descriptivity. Indeed, Edmonds and Moss (2005) claim that descriptive models (in particular agent-based) can help modellers to exploit more of the available data, in particular qualitative information such as anecdotal or commonsense evidence, because of the straightforward correspondence with concepts used in the model; it thus combines well with participatory modelling. Norling also notices that it is easier for experts being modelled to provide information in an intuitive language: 'people have a tendency to explain their actions in terms of their intentions, which in turn are explained in terms of their goals and beliefs' (2004). She thus argues for the use of a descriptive

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

approach such as BDI agents because 'the fact that model builder and subject being modelled are referring to the same concepts does simplify matters'.

Another approach when fine-grained data are not available is to rely on psychological theories of human behaviour to describe expected behaviour, and to use them to design the agent model. For instance, data about emotions is rarely available from field studies or statistics, so theories of emotions can help. In that case, BDI has the advantage of providing the right concepts for capturing these theories and formalise them in terms of mental attitudes. When there is neither data nor theory available, that is, when the decision-making process is unknown, BDI is not an appropriate choice, as also noted in Wolfe *et al*. (2008).

To conclude, BDI agent models might at first seem harder to design because of the lack of fine-grained statistical data, and the unsuitability of such quantitative data to design a qualitative model. But at the same time the BDI formalism allows modellers to exploit different types of data, such as observations from participatory modelling, or psychological theories, because of the straightforward matching between BDI concepts and the way in which people instinctively explain their behaviour.

3.5.2 Computational scalability

Another important criticism against the use of BDI agents in simulation is the computational weight of this approach. As it is already time-consuming to simulate only one BDI agent, with many reasoning components, it is intuitively hard to imagine that such a tool could be used for large-scale simulations with thousands of agents. Some works have integrated BDI agents in real-time environments such as games (see Section 3.3.3) and they tend to confirm this intuition. For example, in Hindriks *et al*. (2010), the authors develop bots using a BDI architecture in the game Unreal Tournament, and conclude that there are indeed problems of scalability and performance when the number of agents increases.

However, this technical drawback is bound to decrease and eventually disappear, as computational power of individual computers increases exponentially, and grid computing also allows the distribution of simulations on many computers to accelerate their execution even more. For instance, Zhao and Son (2007) have built a BDI agent in a distributed environment, to partially replace the decision-making process of human beings in an automated system; this simulation only includes one agent so far but could technically be extended to many more agents. Nevertheless, this appears to be a shift from one research issue to another, because the distribution of a simulator over a network is a very difficult problem too, that is far from being solved.

So scalability still limits the use of BDI agents in large-scale simulations for now, but the computational power of present personal computers already allows the use of BDI agents in smaller scale simulations. For larger scale simulations, when distribution is too complex or limited by a lack of hardware resources, the modeller can resort to an optimised platform, or optimise the implementation of their BDI engine. For instance, Cho *et al*. (2008) implemented a crowd simulator with BDI agents, and evaluated its performance in terms of computing time. They showed that the processing time was roughly proportionate to the number of agents, and mostly dedicated to the belief update process applied every time a new belief is added. Their solution to minimise this computing time was therefore to optimise the reasoner by applying belief update rules only when their effect was significant for the current desire.

In addition, building simulations using BDI agents does not imply that all the agents should be cognitive. A wise choice of which agent should or should not be a BDI agent will have a huge impact on the simulation scalability. Wolfe *et al*. investigate this 'To BDI, or not to BDI' dilemma (2008) in the field of aerial traffic management simulations. They mention that the 'BDI paradigm is an excellent choice when modeling well-understood decision making', but notice efficiency issues when using BDI agents in large-scale simulations. They suggest that one solution is to choose wisely which agents to model as BDI agents (pilots in their example), and which agents could have simpler architectures (planes, etc.). This mixed population of BDI and not-BDI agents increases the performances of their simulator. Novák *et al*. (2013) implemented BDI reasoning capabilities in autonomous military robots teams for surveillance missions, to improve flexibility compared with a simple reactive planner. However, they noted that these agents were heavier and should thus be limited to scenarios where complex deliberation was actually needed.

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

To conclude, BDI architectures are indeed heavier due to the more complex reasoning processes of the agents, so they have to be used wisely. Many authors have successfully applied them in their simulations, by optimising the computation inside each agent, and by limiting their use to agents and scenarios where they were actually justified, mixing BDI agents with simpler agents.

3.5.3 Limited (conceptual) expressivity

As highlighted in Section 2.3.4, the BDI architecture has the advantage of being easily extensible with more concepts than the three basic ones (beliefs, desires, intentions). Nevertheless, before being integrated in a BDI agent, the properties of any additional concept and its links with the existing concepts (at least beliefs, desires and intentions, but possibly others in an already extended BDI architecture) should have been deeply studied.

DSTO (Australian Defence) programmers have listed a number of features that need to be added to the BDI framework so that developers do not have to explicitly implement them in every model, and so that applications can be less constrained. Norling (2004) quotes, for example, emotion, fatigue, learning or memory. We could also extend this list with group attitudes and social commitments. In particular, in a society of agents, the BDI architecture can efficiently be used to represent an individual reasoning process, but only few formalisms have been developed for the reasoning process of social groups. Actually there is no issue metaphorically using a BDI architecture for an agent representing a group of agents when it is well-defined offline (before the beginning of the simulation). But there are only a few studies on how to attribute mental states to social groups of agents emerging online, during the simulation, when the attitudes of the group should interact with the attitudes of the agents composing it.

Only few theoretical studies have been led about the above concepts. There is thus a huge conceptual lack here. Nevertheless, we stay optimistic and assume that in the next years many researchers will try to tackle these issues. For example, recent works have begun to analyze the links between BDI and emotions (Adam, 2007) or social attitudes (Gaudou, 2008). Such work has to be continued to provide a clear, well-grounded and standardised extended BDI framework.

3.5.4 Learning

One of the major *a priori* criticisms of the BDI architecture of agents is its inadequacy to support learning mechanisms: 'one weakness that the BDI agent model shares with other human operator models (agent-based or otherwise) is its inability to support agent learning' (Urlings *et al*., 2006). This is particularly due to the fact that most powerful learning algorithms are dedicated to quantitative learning (learning numerical values, e.g. in neural network, etc.) whereas symbolic learning is a much harder issue (Tables 3).

Nevertheless, there exist some extensions of the BDI architecture aiming at integrating a learning capability into BDI agents. For example, in the video game Black and White (Molyneux, 2001), the main character has a BDI architecture extended with a reinforcement learning capability to allow the player to teach their avatar. In another example, Norling (2004) extended BDI with the RPD strategy used by experts of a domain: the agent progressively learns (from its mistakes) to make very fine distinctions between situations, to allow a nearly automatic plan selection based on it. These examples show that in contrast to preconceived ideas, BDI agents can be useful in simulations even when agents need learning capabilities.

3.6 Conclusion

To conclude this section, we have shown that BDI agents can bring important benefits to social simulation in some cases, depending on the goal of the simulation, the required features of the agents, the desired observation scale, and the field of application. We have also discussed why BDI agents are not used more often in social simulations, and shown that there are solutions to overcome these obstacles.

One more obstacle that we have not discussed yet follows from the high level of abstraction provided by BDI models: this paradigm is a special way of thinking about an autonomous agent, not in terms of classes and methods, but in terms of mental attitudes. As a result, it can be hard for designers to conceptualise their agent model, all the more when they are not programming experts but scientists from other disciplines

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

Table 3 Summary of criticisms against BDI and proposed solutions

	Criticisms	Solutions
Data	Lack of fine-grained quantitative data	Use qualitative data, participatory modelling
	Subjective, non-observable mental states	Straightforward description of psychological theories and expert knowledge
Scalability	Complex, time-consuming reasoning	More powerful computers, grid computing
	Bad performance for real-time or large-scale simulations	
		Optimised reasoning engine
		Mix population of BDI and not-BDI agents
Expressivity	Limited in basic BDI	Extensions exist (norms, emotions, etc.)
	Needs standardised well-grounded extended BDI framework	Recent work towards it
Learning	Most powerful algorithms for quantitative learning, inadequate for BDI (qualitative)	Use other types of learning
		BDI extensions for symbolic learning

(epidemiologists, geographers, sociologists, etc.), as it is often the case in this very multidisciplinary field. Table 3 summarises the criticisms against BDI and our proposed solutions. To facilitate the use of BDI agents in social simulations, it is therefore essential to provide designers with methodologies and tools to support the modelling process and the implementation of the simulator. This is discussed in the next section.

4 Technical guide: how to integrate BDI agents in simulations?

We have seen above that BDI agents bring many benefits to social simulations but that some obstacles impede their wider use. This section tries to provide a guide for concretely integrating BDI agents into social simulations, and to overcome these obstacles.

4.1 Examples of existing simulations with BDI agents

In this article, we have presented a huge array of existing agent-based simulations that use BDI agents. We can first notice that the number of platforms and tools used to implement and run these models is very high. For example, DamasRescue (Paquet *et al*., 2004) models rescue workers after a natural crisis, implements them with Jack agents (Busetta *et al*., 1999), and develops an *ad hoc* simulator to run the simulations. Lui *et al*. (2002) use Jack Teams, an extension of the Jack framework for teams of agents, to model military commanders in land operations scenarios. Cho *et al*. use an existing 3D game engine (SourceEngine) but developed an *ad hoc* BDI reasoner to implement NPC extinguishing a forest fire in a virtual world. Mcilroy *et al*. model fighter pilots by using the dMARS (D'Inverno *et al*., 2004), a realisation of the Procedural Reasoning System, an agent architecture based on the BDI framework. Bosse *et al*. (2007b, 2009) investigate the influence of psychological (emotions) as well as social and biological aspects on criminal behaviour; they use their own language, called LEADSTO, that allows to model both qualitative BDI aspects and quantitative biological aspects.

These few examples are enough to show that many simulations do use BDI agents, but also that most authors have needed to develop at least part of the simulator, be it the BDI engine or the simulation engine. This approach is time-consuming and inefficient, involving the re-development from scratch of *ad hoc* tools instead of reusing existing ones. This could be explained by the lack of awareness of existing tools in other communities, or the greater cost of using them compared with creating *ad hoc* tools.

4.2 BDI platforms that support simulation

BDI platforms provide many useful features (intention selection mechanisms, etc.) that we do not want to have to reprogram in each agent model. However, they miss other necessary features for simulation (e.g. visualisation). For example, Jack Intelligent Agent are the basis of JackSim (Ronald *et al*., 2006) and

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

its extension to teams (Lui *et al*., 2002), which have been used to develop simulators. But all the simulation part was implemented from scratch, or at best coupled with an external visualisation tool.

This example shows that developing a simulator from a BDI platform is not an easy task and requires a lot of work, in particular related to the implementation and management of the environment, the visualisation of the simulation, and the time and scheduling of agents. The management of data is also required for most modern simulators: the simulator should be able to import tabular or geographical data, to create and initialise agents from this data, and to store data in files. Most of these features can now be found in ABMSs platforms such as GAMA (Grignard *et al*., 2013), Netlogo (Wilensky, 1999), Repast Simphony (North *et al*., 2013) or Mason (Luke *et al*., 2005), but are not natively provided by BDI frameworks.

4.3 Simulation platforms that support BDI

In agent-based simulation platforms cited above such as Netlogo, the architecture describing the behaviour of agents is often very simple. Rafael and Jomi (2009), 'shows how trivial the agent behavior is in the type of social simulation that some researchers have had to content themselves with, due to the limitation of agent-based simulation toolkits' (2009). But agent-based simulation toolkits are now evolving, and in particular they sometimes provide more complex architectures, such as the BDI architecture.

For example, Sakellariou have extended the NetLogo platform (Wilensky, 1999) by providing libraries that allow the easy development of BDI agents exchanging FIPA-ACL messages in a FIPA protocol (Sakellariou *et al*., 2008). They developed a demonstration example of a forest fire simulation where firemen agents cooperate and coordinate under the Contract Net Protocol.

Recently, the GAMA platform has also been extended to allow modellers to use a BDI architecture for their agents (Caillou *et al*., 2015). Without being as sophisticated as a dedicated BDI platform, it still provides primitives to manage beliefs, desires and intentions databases in agents and to describe their behaviour in terms of the plans they can perform. The focus of this extension was to keep the implementation simple for modellers, even non-computer scientists.

These examples illustrate a trend to extend simulation platforms with an *ad hoc* BDI module. However, it is not optimal to reprogram BDI reasoning from scratch, while dedicated BDI platforms have been developed and improved over many years. It would therefore be more efficient to connect simulation platforms with existing BDI platforms. For example, Caballero *et al*. (2011) have proposed an integration of JASON into the MASON platform.

Nevertheless, in order to avoid that each simulation platform develops its own connectors to each BDI platform, Padgham *et al*. (2011) have proposed a framework aiming at combining any simulation platform with any BDI platform. Their framework already allows to combine several BDI platforms (e.g. Jason and Jack) with several simulation platforms (e.g. Matsim[4] and Repast). It can be used in a simulation platform to allow agents to delegate their reasoning to a BDI framework. But it can also be used as the master component, driving the simulation platform and synchronising its execution with the BDI platform. It is a very promising tool that could speed up the use of BDI agents in agent-based simulations.

Table 4 considers all the platforms that have been compared in Kravari and Bassiliades (2015) and detail how they support BDI agents.

4.4 Methodologies to help with creating BDI agents

In Agent-Oriented Software Engineering (AOSE), dedicated methodologies and tools exist for expert programmers, such as the Prometheus methodology and associated tool with Jack code generation (Padgham & Winikoff, 2002; Padgham *et al*., 2008); Cossentino *et al*. (2012) describe an UML like notation and Jason-code generation. These methodologies are sometimes used in simulations, for instance, Sokolova and Fernández-Caballero (2007) use Prometheus methodology for modelling agents with a BDI architecture, in a decision-support system for evaluating the impact of pollutant exposure on population health.

[4] www.matsim.org

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

Table 4 Support of agent-based platform (list of platforms from Kravari and Bassiliades (2015)

BDI agents support	ABM platforms
Full native support	Agent Factory, AgentBuilder INGENIAS Development Kit, Jack, Jadex, Jason
Full support with coupling	AgentScape (Jason), Jade (Jadex, Jason), EMERALD (based on Jade), JAMES II MadKit, SeSAm (through Jade)
Limited support	Netlogo (extension), GAMA (plugin) Repast Simphony, Mason
No support, no available information, not relevant	A-Globe, AnyLogic, Cormas, Cougaar, CybelePro JAS, Swarm JIAC (support planned for Spring 2016)

ABM = agent-based modelling.

But while methodologies do exist for BDI agents in AOSE, they are aimed at computer scientists and expert programmers. On the contrary, social simulations are often developed by non-computer scientists, who need appropriate methodologies and tools to help them. In the social simulation field, very few methodologies exist and are used, and existing ones are too high level. For example, the ARDI method (Etienne *et al*., 2011) focusses on the elicitation from stakeholders of a description of the reference system in terms of actors, resources and dynamics; it does not go any deeper in the design of agents and their behaviours. Graphical visualisation tools could also help in involving stakeholders and the general public in participatory modelling (Gaudou *et al*., 2011).

Due to the lack of dedicated methodologies, UML is most often used to conceptualise the agent model. Recently, the Tactics Development Framework (TDF) (Evertsz *et al*., 2014, 2015) was introduced as a methodology for eliciting tactical knowledge from military experts and embedding it in agent-based simulations, and was shown to be better than UML. It was then adapted to modelling the behaviour of the civilian population in Australian bushfires (Adam *et al*., 2015).

To conclude, technologies start to appear to link BDI platforms with modelling and simulation platforms. Their aim is to facilitate the design and implementation by any modellers (especially non-computer scientists) of simulations embedding complex (in particular BDI) agents. Nevertheless, there is still a lack of methodologies to drive the modelling process.

5 Conclusion

We have so far shown the benefits associated with using the BDI architecture for modelling agent behaviour in ABMS. We have talked about the capabilities provided by the core architecture: adaptability, robustness, abstract programming and the ability to explain their behaviour. We have noted that BDI is rarely used in implemented ABMS and we have examined some of the possible reasons for these in turn: the field of application—BDI is unsuitable for modelling simple entities such as bacteria, but ideal for modelling humans; the purpose of the simulation—as recognised by Axelrod, when the simulation requires an accurate representation, complexity is inevitable; the scalability limitations of BDI—these were refuted by real-time game applications; or conceptual lacks of BDI—these can be corrected through extensions such as adding emotions or norms.

Despite the clear benefits, and the lack of good reason against using BDI (at least in some circumstances) there is a lack of uptake of the paradigm in the ABM community. We suggest that this may be because of the lack of appropriate tools to help designers integrate BDI agents into their simulations, and more specifically, the lack of BDI modelling capabilities in existing simulation and modelling platforms. As it is put by Rafael and Jomi: 'the example shows how trivial the agent behavior is in the type of social simulation that some researchers have had to content themselves with, due to the limitation of agent-based simulation toolkits' (2009).

We aim to address this problem in future work by providing tools to facilitate the design of BDI-style agents; in particular we are currently extending GAMA with a BDI architecture, and also working on coupling the TDF methodology with GAMA. We hope that such tools can improve the uptake of BDI agents in social simulations in multiple disciplinary fields of application.

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

Acknowledgements

This work has been partially funded by the ACTEUR ('Spatial Cognitive Agents for Urban Dynamics and Risk Studies') research project. The ACTEUR project is funded by the French National Research Agency (ANR) under grant number ANR-14-CE22-0002.

References

Adam, C. 2007. *The Emotions: From Psychological Theories to Logical Formalisation and Implementation in a BDI Agent*. PhD thesis, IRIT.

Adam, C., Beck, E. & Dugdale, J. 2015. SWIFT: simulations with intelligence for fire training. In *ISCRAM*, Poster.

Adam, C., Gaudou, B., Herzig, A. & Longin, D. 2006. OCC's emotions: a formalization in a BDI logic. In *AIMSA*, Euzenat J. & Domingue J. (eds), LNAI **4183**, 24–32. Springer.

Adam, C., Gaudou, B., Longin, D. & Lorini, E. 2011. Logical modeling of emotions for ambient intelligence. In *Handbook of Research on Ambient Intelligence: Trends and Perspectives*, Mastrogiovanni F. & Chong N.-Y. (eds). IGI Global, 108–127.

Adam, C., Herzig, A. & Longin, D. 2009. A logical formalization of the OCC theory of emotions. *Synthese* **168**(2), 201–248.

Adam, C. & Longin, D. 2007. Endowing emotional agents with coping strategies: from emotions to emotional behaviour. In *7th Intelligent Virtual Agents (IVA)*, 348–349.

Adam, C. & Lorini, E. 2014. A BDI emotional reasoning engine for an artificial companion. In *Highlights of Practical Applications of Heterogeneous Multi-Agent Systems. The PAAMS Collection*, Demazeau, Y., Zambonelli, F., Corchado, J.M. & Bajo Pérez, J. (eds). Springer, 66–78.

Anderson, J. 1983. *The Architecture of Cognition*. Harvard University Press.

Anderson, J. 1993. *Rules of the Mind*. Erlbaum.

Anderson, J. & Lebiere, C. 1998. *The Atomic Components of Thought*. Erlbaum.

Axelrod, R. 1997. *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton University Press.

Axelrod, R. 2005. Advancing the art of simulation in the social sciences. In *Handbook of Research on Nature Inspired Computing for Economy and Management*, Rennard J.-P. (ed.). Idea Group, 21–40.

Axtell, R., Epstein, J., Dean, J., Gumerman, G., Swedlund, A., Harburger, J., Chakravarty, S, Hammond, R., Parker, J. & Parker, M. 2002. Population growth and collapse in a multiagent model of the Kayenta Anasazi in Long House Valley. In *Proceedings of the National Academy of Sciences of the United States of America*.

Balasubramanian, V., Massguer, D. & Mehrotra, S. 2006. Drillsim: a simulation framework for emergency response drills. In *Proceedings of ISCRAM*.

Balke, T. & Gilbert, N. 2014. How do agents make decisions? A survey. *Journal of Artificial Societies and Social Simulation* **17**, `http://jasss.soc.surrey.ac.uk/17/4/13.html`

Baptista, M. L., Martinho, C. R., Lima, F., Santos, P. A. & Prendinger, H. 2014. An agent-based model of consumer behaviour based on the BDI architecture and neoclassical theory. *Developments in Business Simulation and Experiential Learning* **41**, 170–178.

Bazzan, A. L. C., Wahle, J. & Klügl, F. 1999. Agents in traffic modelling: from reactive to social behaviour. In *KI-99: Advances in AI*, Burgard W., Christaller T. & Cremers A. (eds), LNAI **1701**, 303–306. Springer-Verlag.

Bordini, R., Hubner, J. & Wooldridge, M. 2007. *Programming Multi-Agent Systems in AgentSpeak Using Jason*. Wiley-Interscience.

Bosse, T., Gerritsen, C. & Treur, J. 2007a. Cognitive and social simulation of criminal behaviour: the intermittent explosive disorder case. In *AAMAS*.

Bosse, T., Gerritsen, C. & Treur, J. 2007b. Integrating rational choice and subjective biological and psychological factors in criminal behaviour models. In *ICCM'07*, 181–186.

Bosse, T., Gerritsen, C. & Treur, J. 2009. Towards integration of biological, psychological and social aspects in agent-based simulation of violent offenders. *Simulation* **85**(10), 635–660.

Bosse, T., Hoogendoorn, M., Klein, M. C., Treur, J., Van Der Wal, C. N. & Van Wissen, A. 2013. Modelling collective decision making in groups and crowds: integrating social contagion and interacting emotions, beliefs and intentions. *Autonomous Agents and Multi-Agent Systems* **27**(1), 52–84.

Bosse, T., Memon, Z. A. & Treur, J. 2007. A two-level BDI-agent model for theory of mind and its use in social manipulation. In *AISB*.

Bratman, M. 1987. *Intentions, Plans, and Practical Reason*. Harvard University Press.

Broersen, J., Dastani, M., Hulstijn, J., Huang, Z. & van der Torre, L. 2001. The BOID architecture: conflicts between beliefs, obligations, intentions and desires. In *AGENTS'01*, ACM.

Buford, J., Jakobson, G., Lewis, L., Parameswaran, N. & Ray, P. 2006. D-AESOP: a situation aware BDI agent system for disaster situation. In *Agent Technology for Disaster Management*.

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

Busetta, P., Ronnquist, R., Hodgson, A. & Lucas, A. 1999. Jack intelligent agents-components for intelligent agents in java. *AgentLink News Letter* **2**, 2–5.

Caballero, A., Botla, J. & Gomez-Skarmeta, A. 2011. Using cognitive agents in social simulations. *Engineering Applications of Artificial Intelligence* **24**(7), 1098–1109.

Caillou, P., Gaudou, B., Grignard, A., Truong, C. Q. & Taillandier, P. 2015. A simple-to-use BDI architecture for agent-based modeling and simulation. In *11th Conference of the European Social Simulation Association (ESSA)*.

Campenn, M., Andrighetto, G., Cecconi, F. & Conte, R. 2009. Normal = normative? The role of intelligent agents in norm innovation. In *Normative Multi-Agent Systems*.

Castelfranchi, C., Dignum, F., Jonker, C. M. & Treur, J. 2000. Deliberative normative agents: principles and architecture. In *ATAL'99*, LNCS **1757**, 364–378. Springer-Verlag.

Cecconi, F. & Parisi, D. 1998. Individual versus social survival strategies. *Journal of Artificial Societies and Social Simulation* **1**(2), `http://jasss.soc.surrey.ac.uk/1/2/1.html`

Cho, K., Iketani, N., Kikuchi, M., Nishimura, K., Hayashi, H. & Hattori, M. 2008. BDI model-based crowd simulation. In *Intelligent Virtual Agents*, 364–371. Springer.

Cirillo, R., Thimmapuram, P., Veselka, T., Koritarov, V., Conzelmann, G., Macal, C., Boyd, G., North, M., Overbye, T. & Cheng, X. 2006. Evaluating the potential impact of transmission constraints on the operation of a competitive electricity market in Illinois. Report ANL-06/16 for the Illinois Commerce Commission.

Cohen, P. R. & Levesque, H. J. 1990. Intention is choice with commitment. *Artificial Intelligence Journal* **42**(2–3), 213–261.

Conte, R. & Castelfranchi, C. 1995. Understanding the functions of norms in social groups through simulation. In *Artificial Societies: The Computer Simulation of Social Life*, Gilbert N. & Conte R. (eds). Taylor & Francis, 252–267.

Cossentino, M., Chella, A., Lodato, C., Lopes, S., Ribino, P. & Seidita, V. 2012. A notation for modeling Jason-like BDI agents. In *Complex, Intelligent and Software Intensive Systems (CISIS)*, 12–19. IEEE.

Criado, N., Argente, E. & Botti, V. 2010. A BDI architecture for normative decision making. In *AAMAS*, 1383–1384. IFAAMAS.

Criado, N., Argente, E., Noriega, P. & Botti, V. 2010. Towards a normative BDI architecture for norm compliance. In *COIN@MALLOW*, Fornara, N. & Vouros, G. (eds).

DeAngelis, D. & Gross, L. 1992. *Individual-Based Models and Approaches in Ecology*. Chapman and Hall.

Dennett, D. 1989. *The Intentional Stance*. The MIT Press.

de Rosis, F., Pelachaud, C., Poggi, I., Carofiglio, V. & De Carolis, B. 2003. From Greta's mind to her face: modelling the dynamics of affective states in a conversational embodied agent. *International Journal of Human-Computer Studies* **59**(1–2, Special Issue on Application of Affective Computing in HCI), 81–118.

de Silva, L., Sardina, S. & Padgham, L. 2009. First principles planning in BDI systems. In *AAMAS*, 1105–1112.

Dignum, F., Morley, D., Sonenberg, E. & Cavedon, L. 2000. Towards socially sophisticated BDI agents. In *ICCCN*, 0111. IEEE Computer Society.

D'Inverno, M., Luck, M., Georgeff, M., Kinny, D. & Wooldridge, M. 2004. The dMARS architecture: a specification of the distributed multi-agent reasoning system. In *AAMAS*, **9**, 5–53. Kluwer Academic Publishers..

Drogoul, A., Vanbergue, D. & Meurisse, T. 2002. Multi-agent based simulation: where are the agents?. In *MABS*, 1–15.

Edmonds, B. & Moss, S. 2005. From kiss to kids: an anti-simplistic modelling approach. In *MABS*, Davidsson P. (ed.), LNAI **3415**, 130–144. Springer.

Elliott, C., Rickel, J. & Lester, J. 1999. Lifelike pedagogical agents and affective computing: an exploratory synthesis. In *AI Today*, Wooldridge M. & Veloso M. (eds), LNCS **1600**, 195–212. Springer-Verlag.

Etienne, M., Du Toit, D. R. & Pollard, S. 2011. ARDI: a co-construction method for participatory modeling in natural resources management. *Ecology and Society* **16**(1), 44.

Evertsz, R., Thangarajah, J., Yadav, N. & Li, T. 2014. Tactics development framework (demo). In *AAMAS*, 1639–1640.

Evertsz, R., Thangarajah, J., Yadav, N. & Li, T. 2015. Agent oriented modelling of tactical decision making. In *AAMAS*, Bordini, R. H., Elkind, E., Weiss, G. & Yolum, P. (eds), 1051–1060, IFAAMAS.

Fähndrich, J., Ahrndt, S. & Albayrak, S. 2013. Self-explaining agents. *Jurnal Teknologi* **63**(3), 147–154.

Farias, G. P., Dimuro, G. P. & da Rocha Costa, A. C. 2010. BDI agents with fuzzy perception for simulating decision making in environments with imperfect information. In *MALLOW*.

Farmer, J. D. & Foley, D. 2009. The economy needs agent-based modelling. *Nature* **460**, 685–686.

Fernandes, P. & Nunes, U. 2008. Multi-agent architecture for simulation of traffic with communications. In *International Conference on Informatics in Control, Automation and Robotics (ICINCO)*.

Finin, T., Fritzson, R., McKay, D. & Robin, M. 1994. KQML as an agent communication language. In *Proceedings of the Third International Conference on Information and Knowledge Management*.

FIPA 2002a. FIPA communicative act library specification. Foundation for Intelligent Physical Agents. `http://www.fipa.org/specs/fipa00037/`

FIPA 2002b. FIPA contract net interaction protocol specification. Foundation for Intelligent Physical Agents. `http://www.fipa.org/specs/fipa00029/`

Gardner, M. 1970. Mathematical Games. The fantastic combinations of John Conway's new solitaire game 'life'. *Scientific American* **223**, 120–123.

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

Gasmi, N., Grignard, A., Drogoul, A., Gaudou, B., Taillandier, P., Tessier, O. & An, V. D. 2015. Reproducing and exploring past events using agent-based geo-historical models. In *Multi-Agent-Based Simulation XV*, 151–163. Springer.

Gaudou, B. 2008. *Formalizing Social Attitudes in Modal Logic*. PhD thesis, IRIT.

Gaudou, B., Marilleau, N. & Ho, T. V. 2011. Toward a methodology of collaborative modeling and simulation of complex systems. In *Intelligent Networking, Collaborative Systems and Applications*, 27–53. Springer.

Gaudou, B., Sibertin-Blanc, C., Therond, O., Amblard, F., Auda, Y., Arcangeli, J.-P., Balestrat, M., Charron-Moirez, M.-H., Gondet, E., Hong, Y., Lardy, R., Louail, T., Mayor, E., Panzoli, D., Sauvage, S., Sanchez-Pérez, J-M., Taillandier, P., Nguyen, V. B., Vavasseur, M. & Mazzega, P. D. 2013. The MAELIA multi-agent platform for integrated assessment of low-water management issues. In *MABS, Multi-Agent-Based Simulation XIV-International Workshop* (to appear, 2013).

Gilbert, N. & Troitzsch, K. G. 2005. *Simulation for the Social Scientist—Second Edition*. Open University Press.

Gil-Quijano, J., Piron, M. & Drogoul, A. 2007. Mechanisms of automated formation and evolution of social-groups: a multi-agent system to model the intra-urban mobilities of Bogota city. In *Social Simulation : Technologies, Advances and New Discoveries*, Chapter 12, Edmonds B., Hernandez C. & Troitzsch K. (eds). Idea Group Inc., 151–168.

Goldman, A. I. 2012. Theory of mind. In *The Oxford Handbook of Philosophy of Cognitive Science*, Margolis, E., Samuels, R. & Stich, S. P. (eds). Oxford Handbooks Online, `http://www.oxfordhandbooks.com/view/10.1093/oxfordhb/9780195309799.001.0001/oxfordhb-9780195309799-e-17`

Gomboc, D., Solomon, S., Core, M. G., Lane, H. C. & Lent, M. V. 2005. Design recommendations to support automated explanation and tutoring. In *BRIMS05*.

Gratch, J. & Marsella, S. 2004. A domain-independent framework for modeling emotion. *Journal of Cognitive Systems Research* **5**(4), 269–306.

Gratch, J. & Marsella, S. 2005. Some lessons from emotion psychology for the design of lifelike characters. *Journal of Applied Artificial Intelligence* **19**(3–4), Special Issue on Educational Agents -Beyond Virtual Tutors), 215–233.

Gratch, J., Rickel, J., Andre, E., Badler, N., Cassell, J. & Petajan, E. 2002. Creating interactive virtual humans: some assembly required. *IEEE Intelligent Systems* **17**(4), 54–63.

Grignard, A., Taillandier, P., Gaudou, B., Vo, D. A., Huynh, N. Q. & Drogoul, A. 2013. Gama 1.6: advancing the art of complex agent-based modeling and simulation. In *Principles and Practice of Multi-Agent Systems*, 117–131. Springer.

Guiraud, N., Longin, D., Lorini, E., Pesty, S. & Rivière, J. 2011. The face of emotions: a logical formalization of expressive speech acts. In *The 10th International Conference on Autonomous Agents and Multiagent Systems*, **3**, 1031–1038. International Foundation for Autonomous Agents and Multiagent Systems.

Gunderson, L. & Brown, D. 2000. Using a multi-agent model to predict both physical and cyber criminal activity. In *IEEE International Conference on Systems, Man, and Cybernetics* **4**, 2338–2343.

Harbers, M. 2011. *Self-explaining agents*. PhD thesis, Utrecht University.

Harbers, M., van den Bosch, K. & Meyer, J.-J. 2010. Explaining simulations through self-explaining agents. *Journal of Artificial Societies and social simulation* **13**(1), `http://jasss.soc.surrey.ac.uk/13/1/4.html`

Heckbert, S. 2013. MayaSim: an agent-based model of the ancient Maya social-ecological system. *Journal of Artificial Societies and Social Simulation* **16**(4), 11.

Heinze, C., Goss, S., Josefsson, T., Bennett, K., Waugh, S., Lloyd, I., Murray, G. & Oldfield, J. 2001. Interchanging agents and humans in military simulation. In *IAAI*.

Helbing, D., Farkas, I. & Vicsek, T. 2000. Simulating dynamical features of escape panic. *Nature* **407**(6803), 487–490.

Helman, D. H. & Bahuguna, A. 1986. Explanation systems for computer simulations. In *Proceedings of the 18th Conference on Winter Simulation*, 453–459. ACM.

Hindriks, K., De Boer, F., Van der Hoek, W. & Meyer, J. 1999. Agent programming in 3APL. *Autonomous Agents and Multi-Agent Systems* **2**(4), 357–401.

Hindriks, K. V., van Riemsdijk, M. B., Behrens, T., Korstanje, R., Kraaijenbrink, N., Pasman, W. & de Rijk, L. 2010. Unreal goal bots. connecting agents to complex dynamic environments. In *AGS 2010*.

Jones, H., Saunier, J. & Lourdeaux, D. 2009. Personality, emotions and physiology in a BDI agent architecture: the PEP—BDI model. In *Web Intelligence and Intelligent Agent Technologies*.

Karim, S. & Heinze, C. 2005. Experiences with the design and implementation of an agent-based autonomous UAV controller. In *AAMAS*, 19–26. ACM.

Kashif, A., Le, X. H. B., Dugdale, J. & Ploix, S. 2011. Agent based framework to simulate inhabitants' behaviour in domestic settings for energy management. In *ICAART*, 190–199.

Koster, A., Schorlemmer, M. & Sabater-Mir, J. 2012. Opening the black box of trust: reasoning about trust models in a BDI agent. *Journal of Logic and Computation* **23**(1), 25–58.

Kravari, K. & Bassiliades, N. 2015. A survey of agent platforms. *Journal of Artificial Societies and Social Simulation* **18**(1), 11.

Laperriere, V., Badariotti, D., Banos, A. & Muller, J.-P. 2009. Structural validation of an individual-based model for plague epidemics simulation. *Ecological Complexity* **6**(2), 102–112.

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

Lazarus, R. S. 1991. *Emotions and Adaptation*. Oxford University Press.
Lorini, E., Longin, D., Gaudou, B. & Herzig, A. 2009. The logic of acceptance: grounding institutions on agents attitudes. *Journal of Logic and Computation* **19**(6), 901–940.
Lui, F., Connell, R. & Vaughan, J. 2002. An architecture to support autonomous command agents for onesaf testbed simulations. In *SimTecT Conference*.
Luke, S., Cioffi-Revilla, C., Panait, L., Sullivan, K. & Balan, G. 2005. MASON: a multiagent simulation environment. *Simulation* **81**(7), 517–527.
Macal, C. M. & North, M. J. 2005. Tutorial on agent-based modeling and simulation. In *37th Winter Simulation Conference. Introductory Tutorials: Agent-Based Modeling*, 2–15.
Mcilroy, D. & Heinze, C. 1996. Air combat tactics implementation in the smart whole air mission model. In *First International SimTecT Conference*.
Minh, L. V., Adam, C., Canal, R., Gaudou, B., Vinh, H. T. & Taillandier, P. 2012. Simulation of the emotion dynamics in a group of agents in an evacuation situation. In *Principles and Practice of Multi-Agent Systems*, LNCS **7057**, 604–619. Springer.
Molyneux, P. 2001. *Postmortem: Lionhead Studios' Black and White*. Game Developer.
Moss, S., Pahl-Wostl, C. & Downing, T. 2001. Agent-based integrated assessment modelling: the example of climate change. *Integrated Assessment* **2**(1), 17–30.
Murata, S., Arie, H., Ogata, T., Sugano, S. & Tani, J. 2014. Learning to generate proactive and reactive behavior using a dynamic neural network model with time-varying variance prediction mechanism. *Advanced Robotics* **28**(17), 1189–1203.
Nagel, K. & Schreckenberg, M. 1992. A cellular automaton model for freeway traffic. *Journal de physique* **12**(12), 2221–2229.
Nair, R., Tambe, M. & Marsella, S. 2005. The role of emotions in multiagent teamwork. In *Who Needs Emotions: The Brain Meets the Robot*, Fellous J.-M. & Arbib M. (eds), 311–329. Oxford University Press.
Neumann, M. 2010. Norm internalisation in human and artificial intelligence. *Journal of Artificial Societies and Social Simulation* **13**(1), 12.
Newell, A. 1990. *Unified Theories of Cognition*. Harvard University Press.
Norling, E. 2003. Capturing the quake player: using a BDI agent to model human behaviour. In *AAMAS*, 1080–1081.
Norling, E. 2004. Folk psychology for human modeling: extending the BDI paradigm. In *AAMAS*.
Norling, E. J. 2009. *Modelling Human Behaviour with BDI Agents*. PhD thesis, University of Melbourne. `http://cfpm.org/~emma/pubs/thesis.pdf`
Noroozian, A., Hindriks, K. V. & Jonker, C. M. 2014. Towards simulating heterogeneous drivers with cognitive agents. In *ICAART*.
North, M. J., Collier, N. T., Ozik, J., Tatara, E. R., Macal, C. M., Bragen, M. & Sydelko, P. 2013. Complex adaptive systems modeling with repast simphony. *Complex Adaptive Systems Modeling* **1**(1), 1–26.
Novák, P., Komenda, A., Cap, M., Voknnek, J. & Pechoucek, M. 2013. Simulated multi-robot tactical missions in urban warfare. In *Multiagent Systems and Applications*, 147–183. Springer.
Ortony, A., Clore, G. & Collins, A. 1988. *The Cognitive Structure of Emotions*. Cambridge University Press.
Ostrom, E. 2007. A general framework for analyzing sustainability of social-ecological systems. *Proceedings of the Royal Society of London, Series B* **274**, 1931.
Oulhaci, M. A., Tranvouez, E., Fournier, S. & Espinasse, B. 2013. A multi-agent architecture for collaborative serious game applied to crisis management training: Improving adaptability of non played characters. In *European Conference on Games Based Learning*, 465.
Padgham, L., Scerri, D., Jayatilleke, G. & Hickmott, S. 2011. Integrating BDI reasoning into agent based modeling and simulation. In *Proceedings of the Winter Simulation Conference*, 345–356.
Padgham, L., Thangarajah, J. & Winikoff, M. 2008. Prometheus design tool. In *23rd AAAI Conference on AI*, 1882–1883. AAAI Press.
Padgham, L. & Winikoff, M. 2002. Prometheus: a methodology for developing intelligent agents. In *AOSE @ AAMAS*.
Palazzo, L., Dolcini, G., Claudi, A., Biancucci, G., Sernani, P., Ippoliti, L., Salladini, L. & Dragoni, A. F. 2013. Spyke3d: a new computer games oriented BDI agent framework. In *2013 18th International Conference on Computer Games: AI, Animation, Mobile, Interactive Multimedia, Educational & Serious Games (CGAMES)*, 49–53. IEEE.
Paquet, S., Bernier, N. & Chaib-draa, B. 2004. DAMAS-rescue description paper. In *Proceedings of RoboCup-2004: Robot Soccer World Cup VIII*, 12. Springer-Verlag.
Park, S. I. 2013. *Modeling Social Group Interactions For Realistic Crowd Behaviors*. PhD thesis, Virginia Polytechnic Institute and State University.
Peinado, F., Cavazza, M. & Pizzi, D. 2008. Revisiting character-based affective storytelling under a narrative BDI framework. In *Interactive Storytelling*, 83–88. Springer.
Pereira, D., Oliveira, E., Moreira, N. & Sarmento, L. 2005. Towards an architecture for emotional BDI agents. In *Twelfth Portuguese Conference on AI*, 40–47.
Pokahr, A., Braubach, L. & Lamersdorf, W. 2005. Jadex: a BDI reasoning engine. In *Multi-Agent Programming*, 149–174. Springer.

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

Posada, M., Hernaandex, C. & Laopez-Paredes, A. 2008. Emissions permits auctions: an agent based model analysis. In *Social Simulation: Technologies, Advances and New Discoveries*, Chapter 14 Edmonds B., Troitzsch K. G. & Iglesias C. H. (eds). IGI Global, 180–191.

Rafael, H. B. & Jomi, F. H. 2009. Agent-Based Simulation Using BDI Programming in Jason. In *Multi-Agent Systems: Simulation and Applications*, Uhrmacher, A. M. & Weyns, D. (eds), 451–476. Taylor and Francis Group.

Rao, A. S. & Georgeff, M. P. 1991. Modeling rational agents within a BDI-architecture. In *KR'91*, Allen J. A., Fikes R. & Sandewall E. (eds). Morgan Kaufmann, 473–484.

Reynolds, C. 1987. Flocks, herds, and schools: a distributed behavior model. In *SIGGRAPH*.

Rickel, J., Gratch, J., Hill, R., Marsella, S. & Swartout, W. 2001. Steve goes to Bosnia: towards a new generation of virtual humans for interactive experiences. In *AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment*, Stanford University.

Rivière, J., Adam, C. & Pesty, S. 2012. A reasoning module to select ECAs communicative intention. In *Intelligent Virtual Agents*, 447–454. Springer.

Rizzo, A., Kenny, P. & Parsons, T. D. 2011. Intelligent virtual patients for training clinical skills. *Journal of Virtual Reality & Broadcasting* **8**(3), `https://www.jvrb.org/past-issues/8.2011/2902`

Ronald, N., Sterling, L. & Kirley, M. 2006. Evaluating Jack Sim for agent-based modelling of pedestrians. In *Intelligent Agent Technology (IAT)*, 81–87. IEEE/WIC, ACM.

Rönnquist, R. 2008. The goal oriented teams (Gorite) framework. In *Programming Multi-Agent Systems*, 27–41. Springer.

Rosenbloom, P. S., Laird, J. E. & Newell, A. 1993. *The SOAR Papers: Research on Integrated Intelligence*. MIT Press.

Sakellariou, I., Kefalas, P. & Stamatopoulou, I. 2008. Enhancing netlogo to simulate BDI communicating agents. *AI: Theories, Models and Applications* **5138**, 263–275.

Sardina, S., de Silva, L. & Padgham, L. 2006. Hierarchical planning in BDI agent programming languages: a formal approach. In *AAMAS'06*, 1001–1008. ACM.

Savarimuthu, B. T. R. & Cranefield, S. 2009. A categorization of simulation works on norms. In *Dagstuhl Seminar Proceedings 09121: Normative Multi-Agent Systems*, 39–58.

Schattenberg, B. & Uhrmacher, A. 2001. Planning agents in James. *Proceedings of the IEEE* **89**(2), 158–173.

Schelling, T. C. 1971. Dynamic models of segregation. *Journal of Mathematical Sociology* **1**, 143–186.

Shendarkar, A., Vasudevan, K., Lee, S. & Son, Y.-J. 2006. Crowd simulation for emergency response using BDI agents based on immersive virtual reality. In *Winter Simulation Conference*, L. F. Perrone et al. (ed.), 545–553. ACM.

Shendarkar, A., Vasudevan, K., Lee, S. & Son, Y.-J. 2008. Crowd simulation for emergency response using BDI agents based on immersive virtual reality. *Simulation Modelling Practice and Theory* **16**(9), 1415–1429.

Silverman, B. G., Badler, N. I., Pelechano, N. & O'Brien, K. 2005. Crowd simulation incorporating agent psychological models, roles and communication. Retrieved from `http://repository.upenn.edu/hms/29`

Simoes, J. A. 2012. An agent-based/network approach to spatial epidemics. In *Agent-Based Models of Geographical Systems*, 591–610. Springer.

Singh, D., Sardina, S., Padgham, L. & James, G. 2011. Integrating learning into a BDI agent for environments with changing dynamics. In *22nd IJCAI*.

Small, R. K. 2008. Agent smith: a real-time game-playing agent for interactive dynamic games. In *Genetic and Evolutionary Computation Conference*, 1839–1842.

Sokolova, M. V. & Fernández-Caballero, A. 2007. An agent-based decision support system for ecological-medical situation analysis. In *Nature Inspired Problem-Solving Methods in Knowledge Engineering*, 511–520. Springer.

Staller, A. & Petta, P. 2001. Introducing emotions into the computational study of social norms: a first evaluation. *Journal of Artificial Societies and Social Simulation* **4**(1), `http://jasss.soc.surrey.ac.uk/4/1/2.html`

Sterman, J. D. 2006. Learning from evidence in a complex world. *Journal of Public Health* **96**(3), 505–514.

Steunebrink, B., Dastani, M. & Meyer, J.-J. 2008. A formal model of emotions: integrating qualitative and quantitative aspects. In *ECAI'08*, 256–260. IOS Press.

Sun, R. 2002. *Duality of the Mind*. Lawrence Erlbaum Associates.

Sun, R. 2004. Desiderata for cognitive architectures. *Philosophical Psychology* **17**(3), 341–373.

Sun, R. (ed.) 2006. *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation*. Cambridge University Press.

Sun, R. 2007. Cognitive social simulation incorporating cognitive architectures. *IEEE Intelligent Systems* **22**(5), 33–39.

Sun, R., Merrill, E. & Peterson, T. 2001. From implicit skills to explicit knowledge: a bottom-up model of skill learning. *Cognitive Science* **25**(2), 203–244.

Sun, R. & Naveh, I. 2007. Social institution, cognition, and survival: a cognitive-social simulation. *Mind & Society* **6**(2), 115–142.

Swartout, W., Gratch, J., Hill, R., Hovy, E., Marsella, S., Rickel, J. & Traum, D. 2006. Toward virtual humans. *AI Magazine* **27**(2), 96–108.

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press

Swartout, W., Paris, C. & Moore, J. 1991. Explanations in knowledge systems: design for explainable expert systems. *IEEE Expert* **6**(3), 58–64.

Swarup, S., Eubank, S. G. & Marathe, M. V. 2014. Computational epidemiology as a challenge domain for multiagent systems. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems*, 1173–1176. International Foundation for Autonomous Agents and Multiagent Systems.

Taibi, T. 2010. Incorporating trust into the BDI architecture. *International Journal of Artificial Intelligence and Soft Computing* **2**(3), 223–230.

Taillandier, P., Thrond, O. & Gaudou, B. 2012. A new BDI agent architecture based on the belief theory. Application to the modelling of cropping plan decision-making. In *International Environmental Modelling and Software Society (iEMSs)*, Seppelt R., Voinov A. A., Lange S. & Bankamp D. (eds), 2463–2470. International Environmental Modelling and Software Society.

Tesfatsion, L. 2002. Agent-based computational economics: growing economies from the bottom up. *Artificial Life* **8**, 55–82.

Thabet, I., Hanachi, C. & Ghaedira, K. 2009. Towards an adaptive grid scheduling: architecture and protocols specification. In *Agent and Multi-Agent Systems: Technologies and Applications*, 599–608. Springer.

Thierry, H., Vialatte, A., Choisis, J.-P., Gaudou, B. & Monteil, C. 2014. Managing agricultural landscapes for favouring ecosystem services provided by biodiversity: a spatially explicit model of crop rotations in the GAMA simulation platform. In *International Environmental Modelling and Software Society (iEMSs)*, Ames D. P. & Quinn N. (eds). `http://www.iemss.org/sites/iemss2014/papers/iemss2014_submission_37.pdf`

Traum, D., Rickel, J., Gratch, J. & Marsella, S. 2003. Negotiation over tasks in hybrid human-agent teams for simulation-based training. In *2nd International Conference on Autonomous Agents and Multiagent Systems*.

Traum, D., Swartout, W., Marsella, S. & Gratch, J. 2005. Fight, flight, or negotiate: believable strategies for conversing under crisis In *5th International Conference on Interactive Virtual Agents*.

Urlings, P., Sioutis, C., Tweedale, J., Ichalkaranje, N. & Jain, L. 2006. A future framework for interfacing BDI agents in a real-time teaming environment. *Journal of Network and Computer Applications* **29**(2-Innovations in agent collaboration), 105–123.

Van Truong, H., Beck, E., Dugdale, J. & Adam, C. 2013. Developing a model of evacuation after an earthquake in Lebanon. In *ISCRAM-Vietnam*.

von Wright, G. H. 1963. *Norm and Action*. Routledge and Kegan.

Wan, Y., Zhang, D.-y. & Jiang, Z.-h. 2013. Decision-making algorithm of an agent's internal behavior facing artificial market. *Soft Computing with Applications(SCA)* **1**(1), 20–27.

Wilensky, U. 1999. NetLogo. `http://ccl.northwestern.edu/netlogo/`

Wilks, Y. (ed.) 2010. *Close Engagements With Artificial Companions: Key Social, Psychological, Ethical and Design Issues, Natural Language Processing*. John Benjamins Pub Co.

Wolfe, S. R., Sierhuis, M. & Jarvis, P. A. 2008. To BDI, or not to BDI: design choices in an agent-based traffic flow management simulation. In *Spring Simulation Multiconference*, 63–70. International Society for Computer Simulation.

Wooldridge, M. 2009. *An Introduction to MultiAgent Systems - Second Edition*. John Wiley & Sons.

Ye, L. R. & Johnson, P. E. 1995. The impact of explanation facilities on user acceptance of expert systems advice. *MIS Quaterly* **19**(2), 157–172.

Zhao, X. & Son, Y.-J. 2007. BDI-based human decision-making model in automated manufacturing systems. *International Journal of Modeling and Simulation* **28**(3), 347–356.

https://doi.org/10.1017/S0269888916000096 Published online by Cambridge University Press