

Adaptable and Verifiable BDI Reasoning*

Peter Stringer Rafael C. Cardoso Xiaowei Huang Louise A. Dennis

University of Liverpool
Liverpool L69 3BX, United Kingdom

{peter.stringer, rafael.cardoso, xiaowei.huang, l.a.dennis}@liverpool.ac.uk

Long-term autonomy requires autonomous systems to adapt as their capabilities no longer perform as expected. To achieve this, a system must first be capable of detecting such changes. In this position paper, we describe a system architecture for Belief-Desire-Intention (BDI) autonomous agents capable of adapting to changes in a dynamic environment and outline the required research. Specifically, we describe an agent-maintained self-model with accompanying theories of durative actions and learning new action descriptions in BDI systems.

1 Introduction

Long-term autonomy requires autonomous systems to adapt as their capabilities no longer perform as expected. To achieve this, a system must first be capable of detecting such changes. Creating and maintaining a *system ontology* is a comprehensive solution for this; an agent-maintained formal *self-model* will take the role of this system ontology. It would act as a repository of information about all the processes and functionality of the autonomous system, forming a systematic approach for detecting action failures.

Our work will focus on Belief-Desire-Intention (BDI) [25] programming languages as they are well known for their use in developing intelligent agents [1, 6, 16, 21]. Agents that are capable of controlling an array of cyber-physical autonomous systems such as autonomous vehicles, spacecraft and robot arms have been programmed using BDI agents (e.g., Mars Rover [16], Earth-orbiting satellites [6] and robotic arms for nuclear waste-processing [1]). Coupled with their use of plans and actions, BDI languages offer an appropriate platform to build upon for the development of an adaptable autonomous system.

The agent-maintained self-model will include *action descriptions*, consisting of pre- and post-conditions of all known actions/capabilities. An action's pre-conditions are the environment conditions that must exist for an action to be executed whilst post-conditions are defined as the expected changes in the environment made directly by a completed action. These action descriptions are based on the Planning Domain Definition Language (PDDL) [22], commonly used in classical automated planning. The complete availability of current system information will provide the ability to monitor the status of actions, presenting the opportunity to detect failure. We use *action life-cycles* based on a theory of durative actions for BDI systems [10] to detect persistent abnormal behaviour from action executions that could denote hardware degradation or other long-term causes of failure such as exposure to radiation or extreme temperature. Once a failure has been detected, we can use machine learning methods to update the action description in the self model. Then, we can repair or replace the actions in any existing plans by using an automated planner to patch these plans. The resulting plans can then be verified to ensure the system's safety properties are intact.

*This work has been supported by the University of Liverpool's School of EEE & CS in support of the EPSRC "Robotics and AI for Nuclear" (EP/R026084/1) and "Future AI and Robotics for Space" (EP/R026092/1) Hubs.

This is a position paper that outlines a program of research. The overarching aim of this research is to create a framework for the verification of autonomous systems that are capable of learning new behaviour(s) descriptions and integrating them into existing BDI plans: using the framework as a route to certification. In this paper, we discuss the current ability of BDI systems in adaptable reasoning, largely focusing on actions. We also consider research in Artificial Intelligence (AI) planning on modelling actions and the methods and implications of introducing machine learning for replacing action descriptions. Our main contribution is the initial design of a system architecture for BDI autonomous agents capable of adapting to changes in a dynamic environment, consolidating the agent-maintained self-model with the theory of durative actions and learning new action descriptions into a cohesive and adaptable BDI system. It should be noted, our work relies upon assumptions that have been discussed further in the relevant sections of this paper.

2 The Belief-Desire-Intention Model of Agency

Intelligent Agent systems conforming to the BDI software model largely follow the principles proposed by Bratman's *Intentions, Plans, and Practical Reason* [4] that was originally intended for modelling practical reasoning in human psychology. Use of BDI agents is particularly suitable for high-level management and control tasks in complex dynamic environments [26] which justifies their implementation in many practical applications. Since Georgeff and Lansky's Procedural Reasoning System (PRS) emerged in 1987 [14] a wide range of BDI programming languages have been developed [21], each building upon the reach of PRS with a multitude of extensions for different applications.

2.1 Actions in BDI Systems

Typically, BDI languages use *plans* provided by a programmer at compile time and the language selects an appropriate plan to react to a given situation. Some BDI languages model interaction with the external environment either as an *action* (e.g., Jason [3]) or as a *capability* (e.g., GOAL [19]). We view capabilities as actions with explicit pre- and post-conditions.

Actions and capabilities can appear in the bodies of *plans*. The body of a plan is generally a sequence of actions/capabilities, belief updates and subgoal manipulations (e.g., adopting or dropping goals). Plans are selected by means-end reasoning in order to achieve the agent's goals. Plans may have additional components as well as the plan body. For instance, they generally have a *guard* which must hold before the plan can be applied. Once a plan is selected for execution it is transformed into an *intention* which represents a sequence of steps to be performed as part of executing the plan.

We intend to extend the GWENDOLEN agent programming language [7] for our research. We have chosen to use GWENDOLEN as it is a BDI agent programming language capable of producing verifiable agents. It is integrated into the MCAPL (Model-Checking Agent Programming Languages) framework [8]. Using MCAPL, agents can be programmed in GWENDOLEN and then verified using the AJPF (Agent Java Pathfinder) model-checker [11]. Actions in GWENDOLEN are generally implemented in a Java-based environment; at runtime they are requested and executed by agents. Whilst actions in GWENDOLEN can exhibit characteristics of duration, this is implemented using a 'wait for' construct which temporarily suspends an intention when encountered. When the predicate that is being waited for is believed, the intention becomes unsuspending. This is the extent to which actions in GWENDOLEN are treated as having significant durations and is largely typical of the treatment of actions in BDI languages.

2.2 Durative Actions

BDI languages are increasingly being used for developing agents for physical systems where actions could take considerable time to complete [10]. Currently, most BDI languages suspend an agent entirely until an action completes or implement actions in such a way that an agent may start a process but then must be programmed to explicitly track the progress of the action in some way.

Introducing an explicit notion of duration to actions will allow us to create principled mechanisms to let an agent continue operating once an action is started, meaning the agent is available to monitor the status of actions in progress. [16] introduced an abstract theory of *goal life-cycles*, whereby every goal pursued by the agent moves through a series of states: *Pending* to *Active*; *Active* to either *Suspended* or *Aborted* or a *Successful* end state; and so on. Dennis and Fisher [10] extended the formal semantics provided by Harland et al. to show how the behaviour of durative actions could integrate into these life-cycles. They advocate associating actions not only with pre- and post-conditions containing durations but also with explicit success, failure and abort conditions (an abort is used if the action is ongoing but needs to be stopped) and suggest goals be suspended while an action is executing and then the action's behaviour be monitored for the occurrence of its success, failure or abort conditions. When one of these occurs the goal then moves to the *Active* or *Pending* (where re-planning may be required) part of its life-cycle as appropriate. Adding these additional states to actions should not add to the cost of model checking as this should not add branches. Adding states should only add more information which would make no significant difference.

Brahms [28], a multi-agent modelling environment, is an example of an agent approach that implements durative actions. The Brahms equivalent of actions, *activities*, have duration. Brahms has a formal semantics provided by Stocker et al. [29], although these semantics are primarily concerned with the effect of activity duration on simulation with mechanism for monitoring the behaviour of an activity during its execution. Whilst the concept of durative actions seems to have been adequately explored in these examples, there has not been a formal implementation that focuses on monitoring individual actions for failure.

2.3 Action Failure

The idea of monitoring an action's life-cycle exists in current literature [10, 16, 17]. A range of states can be attributed to an action that can subsequently be traced for irregularities or consistent errors, providing a basis for determining failure. If we assume that the performance of actions may degrade then we also need to introduce the concept of an action life-cycle in which an action is introduced into the system as *Functional*, may move into a *Suspect* state if it is failing and finally becomes *Deprecated* following repeated failures.

Cardoso et al. [5] assumes a framework along these lines and builds upon it to outline a mechanism that allows reconfiguration of the agent's plans in order to continue functioning as intended if some action has become *Deprecated*. However, this assumed ability to detect persistent failures does not yet exist. Our proposed framework should allow us to detect persistent abnormal behaviour from action executions for use with Cardoso et al.'s reconfiguration mechanism.

3 AI Planners and Learning New Actions

AI Planning seeks to automate reasoning about plans; using a formal description of the domain, all possible actions available in the domain, an initial state of the problem, and a goal condition to produce

a plan consisting of the actions that will achieve the goal condition when executed [18]. The formal description of the domain and the problem can be considered a model of the environment, the accuracy of which is fundamental to producing viable plans of reasonable quality. Significant advances have been made in the modelling of actions [12, 15, 33, 34] in automated planning, supporting actions that can have variable duration, conditions and effects.

Actions in BDI systems are typically designed without a specified duration and are defined before the execution of the program. As previously mentioned, BDI systems do not have a de facto theory of durative actions. Additionally, there is no theory for learning new action descriptions. With an extension of action theory for BDI systems (covering these two areas) paired with the self-model concept, actions could adapt to change. However, learning a new action may not always be the best solution for a failing action. Cardoso et al. [5] have developed a method for reasoning about replacing malfunctioning actions with alternate existing actions to achieve the same desired goal, reusing the domain entities and predicates that are already available.

In situations where a new action description is required at runtime, there are already suitable learning methods that could be adapted to be incorporated into the framework [23, 24], enabling the discovery of new entities and predicates in the domain. In [23], the Qualitative Learner of Action and Perception (QLAP) is introduced. When deployed in an unknown, continuous and dynamic environment, QLAP constructs a hierarchical structure of possible actions in an environment based upon the consequences of actions that have happened before. Work in [24] explores the use of Machine Learning and probabilistic planning in complex environments to cope with unexpected outcomes. A learning algorithm is used to determine an action model with the greatest likelihood of attaining the perceived action effects of another different set of actions. We have to acknowledge the risk that system properties could be violated during machine learning, although this could be remedied by using a *Safe Learning* [13] approach, the introduction of machine-learning presents great difficulty for verification as the algorithms cannot (currently) be directly verified [31]. As a consequence, it should be noted that the proposed system could be unsuitable for scenarios where learning from failure is not safe (e.g., autonomous drones), where it would be safest to execute a controlled stop to the system rather than attempting a recovery.

4 System Architecture

The initial objective of this research is to formally define the concept of a self-model: an agent-maintained ontology for the autonomous system. We intend to use PDDL (Planning Domain Definition Language) as a starting point for creating a self-model. PDDL is a formalism for AI planning which is intended to “express the ‘physics’ of a domain” [22]. More specifically, we intend to use features introduced in PDDL2.1 [12] as our starting point. PDDL2.1 is an extension to PDDL for expressing temporal planning domains. The self-model concept will build on this by enabling agents to access and maintain a domain description, adding the capability of learning new action descriptions and allowing action life-cycles to be monitored. As shown in Figure 1, the self-model is centrally linked to the other system components as they are required to contribute into keeping the self-model accurate. It is important to note that the self-model’s domain description is not assumed to be modelled soundly and completely, yet it is assumed that all reports and updates received by the system are correct.

Our implementation will be developed for GWENDOLEN [9]. The GWENDOLEN agent programming language follows the BDI software model. As part of the MCAPL Framework, GWENDOLEN interfaces with the Java Pathfinder (JPF) model-checker [32]. Our intention is to implement self-models and the theory of action life-cycles [10] in GWENDOLEN and integrate this with the existing work on plan re-

configurability [5]. We will then exploit GWENDOLEN’s support for verification to verify the adapted system against requirements. We propose representing actions in our self-model with explicit pre- and post-conditions and either explicit success, fail and abort conditions or one’s that can be inferred from the pre- and post-conditions. We will then adapt the GWENDOLEN goal life-cycle as suggested in [10] to handle durative actions in a principled fashion.

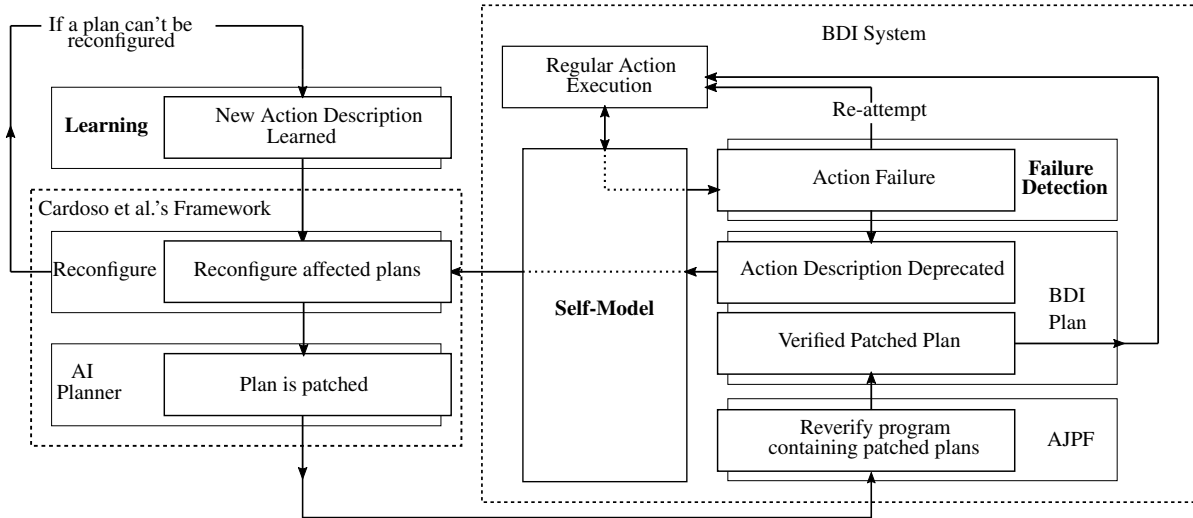


Figure 1: Diagram of Action Failure and Recovery Mechanisms. Arrows represent data flow and dotted lines are for readability when a line goes through a component.

When an action changes, requiring plans to be modified, it is assumed that the agent must be verified again in order to preserve the safety properties of the system as a whole. However, if a new action is learnt in place of a failing action (fully or partially achieving the failing action’s post-conditions), the whole system may not require reverification. We aim to further study this process in order to identify the conditions where such reverification would not be necessary.

In *Algorithm 1*, we propose a primitive method for action failure monitoring. It is assumed that the action status used in the algorithm is asserted as a belief by the system. We start monitoring an action once it has been executed, retrieving some preliminary information about the action: an identifier and the current status (lines 2-3). If the action is currently *Pending*; *Suspended*; or *Aborting*, this status is returned (lines 4-5). If not, the action’s expected post-conditions are retrieved from the tuple (line 6). Whilst an action’s state is *Active*, we continue checking it for failure by comparing the perceived post-conditions with those that are expected of that action (lines 7-11). If at any point during monitoring these conditions do not match, the action’s state becomes *Failed*. If an action is not working as expected, the action can be re-attempted or suspended and replaced in the self-model. The replacement action may be selected from an existing action in the self-model itself. Alternatively, using machine-learning methods, a new action can be learnt to replace the failing action using current knowledge of the available capabilities. Finally, a method for reconfiguring the BDI plan, such as in [5], is called (lines 12-15).

4.1 Scenario

To illustrate how the self-model would complement Cardoso et al.’s work on reconfigurability [5], we use the same scenario: a Mars rover’s faulty movement capability. Figure 1 shows our proposed mechanisms

Algorithm 1: Action Failure Monitoring

```

1 Function monitor( $\langle action\_identifier, action\_status, action\_post\_conditions \rangle$ )
2    $ActionID \leftarrow action\_identifier$ ;
3    $Status \leftarrow action\_status$ ;
4   if  $Status \neq Active \vee Failed$  then
5     return  $Status$ ;
6    $ExpectedPostCond \leftarrow action\_post\_conditions$ ;
7   while  $Status = Active$  do
8      $ActualPostCond \leftarrow getPostConditions(ActionID)$ ;
9     if  $ActualPostCond = ExpectedPostCond$  then
10       $Status \leftarrow Active$ ;
11      monitor( $ActionID, Status, ActualPostCond$ );
12    else
13       $Status \leftarrow Failed$ ;
14      reconfigure( $ActionID$ );
15      return  $Status$ ;
16  return  $Active$ ;

```

for actions failure and recovery embedded into a system architecture including a BDI system, an AI Planner and Cardoso et al.'s [5] reconfigurability framework. The dotted line arrows crossing the self-model represent incoming information from a component such as an action's state. The system architecture in the diagram relies upon a simplification of the successful, fault-free execution of actions that would normally occur in a BDI system. In the case of the rover, these fault-free actions can be represented by the high-level task of movement between waypoints. Whilst mostly successful, these actions are susceptible to failure.

Consider the task of moving from a waypoint A to another waypoint B, in order to collect a rock sample to analyse at another waypoint C. Using the monitoring method for failure detection in Algorithm 1, a failed movement action between any of these waypoints could be found. Given the dynamic environment that the rover operates in, it is plausible that previously clear and usable routes could become blocked at any time. A failure can be flagged when an action is exceeding a predetermined time or energy threshold described in the action post-conditions. Once failure has been detected and confirmed, we can update the self-model to show that the action description has deprecated and no longer affords its post-conditions. The rover now attempts to reconfigure the current plan to resolve the failure using an AI planner to search for a replacement (e.g., by finding a different route) before attempting to learn a completely new action description. In both cases, the time and energy consumption required to accomplish the original post-conditions is updated in the reconfigured/new action description. If it is found that the reconfigured plan is now too time or energy intensive, the latter method of learning a new action description is invoked. If at any point the failing action is found to achieve all post-conditions but does not perform the action within the time or energy threshold (e.g., the rover now navigates around a blockage and arrives at the correct waypoint but now takes longer to do so), this can be managed by learning new actions descriptions with an updated time and/or energy threshold.

The action may not be deprecated if the failure is considered anomalous; for instance, the action normally succeeds and only fails on one isolated occasion. If the action description is not deprecated,

then the action will be re-attempted without resorting to reconfiguration or learning methods. The rover can continue progressing towards the goal if the failure was anomalous. If a new action description is learned, the original plan will be patched with the new action description by the AI planner. This plan could require reverification to preserve previously verified properties which can be handled by using AJPF. Ensuring these properties are maintained is crucial for avoiding failure. The verified patched plan can then be used in the BDI system where regular action execution continues and the rover can continue to complete its mission.

5 Related Work

The work in [5] describes a reconfigurability framework that is capable of replacing faulty action descriptions based on formal definitions of action descriptions, plans, and plan replacement. The implementation uses an AI planner to search for viable action replacements. We plan on extending their approach by adding the concept of a self-model, durative actions, and failure detection. Furthermore, we also envision adding a learning component to the framework in order to be able to cope with dynamic environment events that require new action descriptions to be formulated at runtime.

Troquard et al.'s work on logic for agency in [30] considers the modelling of actions with durations although a different approach was taken: actions are given duration using continuations from STIT (Seeing To It That) logic. In BDI systems, the focus of handling plan failure is the effect that failure has on goals [2, 27]. This is a reasonable focus considering the central role that goals have in agent-oriented programming. Consequently, action failure recovery has not been explored as an option for managing plan failure.

6 Conclusions and Future Work

In this position paper we have described a system architecture for BDI autonomous agents capable of adapting to changes in a dynamic environment. We also introduced the idea of an agent-maintained self-model with durative actions and learning new action descriptions. Our proposed system aims to resolve the following: develop the concept of a self-model; produce and develop a method to detect the failure of an action performed by a BDI Agent; develop a theory of durative actions for BDI languages; adapt existing system to allow new actions to be learnt and used in place of failing ones whilst preserving safety properties, and finally to integrate into the existing GWENDOLEN infrastructure.

To illustrate the applicability of the discussed mechanisms, a practical example of how a Mars rover could make use of the framework was provided. Future work includes defining the learning component to be able to handle dynamic environment events that require the creation of new action descriptions at runtime, a formal definition of the self-model with an outline of the concepts included in this, the implementation of the system architecture, and the evaluation of the approach.

A number of questions and challenges have been identified whilst outlining this program of research. Firstly, it has been noted that the term 'persistent failure' is subjective and should be accompanied by a formal and precise specification to avoid ambiguity. Secondly, considerations for the steps taken after reconfiguration and the learning process require further work (e.g. What happens to failing actions in the model after reconfiguring?). Finally, the proposed learning strategy has produced many challenges which will be considered once implementation has reached this stage. Notably, we will consider how the learning method can ensure valid solutions; how planning time could be minimised and how an action's state could influence the learning strategy. These challenges will serve as guidance for future work.

References

- [1] Jonathan M. Aitken, Affan Shaukat, Elisa Cucco, Louise A. Dennis, Sandor M. Veres, Yang Gao, Michael Fisher, Jeffrey A. Kuo, Thomas Robinson & Paul E. Mort (2017): *Autonomous Nuclear Waste Management*. *IEEE Intelligent Systems*. In Press.
- [2] Rafael H Bordini & Jomi Fred Hübner (2010): *Semantics for the Jason Variant of AgentSpeak (Plan Failure and some Internal Actions)*. In: *ECAI*, pp. 635–640, doi:10.3233/978-1-60750-606-5-635. Available at <https://dx.doi.org/10.3233/978-1-60750-606-5-635>.
- [3] Rafael H. Bordini, Michael Wooldridge & Jomi Fred Hübner (2007, 273p): *Programming Multi-Agent Systems in AgentSpeak using Jason*. John Wiley & Sons, doi:10.1002/9780470061848.
- [4] Michael Bratman et al. (1987): *Intention, plans, and practical reason*. 10, Harvard University Press Cambridge, MA.
- [5] Rafael C. Cardoso, Louise A. Dennis & Michael Fisher (2019): *Plan Library Reconfigurability in BDI Agents*. In: *Proc. of the 7th International Workshop on Engineering Multi-Agent Systems (EMAS)*.
- [6] Louise Dennis, Michael Fisher, Alexei Lisitsa, Nicholas Lincoln & Sandor Veres (2010): *Satellite control using rational agent programming*. *IEEE Intelligent Systems* 25(3), pp. 92–97, doi:10.1109/mis.2010.88.
- [7] Louise A. Dennis (2017): *Gwendolen Semantics: 2017*. Technical Report ULCS-17-001, University of Liverpool, Department of Computer Science.
- [8] Louise A Dennis (2018): *The MCAPL Framework including the Agent Infrastructure Layer and Agent Java Pathfinder*. *The Journal of Open Source Software*.
- [9] Louise A Dennis & Berndt Farwer (2008): *Gwendolen: a BDI language for verifiable agents*. In: *Proceedings of the AISB 2008 Symposium on Logic and the Simulation of Interaction and Reasoning*, Society for the Study of Artificial Intelligence and Simulation of Behaviour, pp. 16–23.
- [10] Louise A Dennis & Michael Fisher (2014): *Actions with Durations and Failures in BDI Languages*. In: *ECAI*, pp. 995–996, doi:10.3233/978-1-61499-419-0-995. Available at <https://dx.doi.org/10.3233/978-1-61499-419-0-995>.
- [11] Louise A. Dennis, Michael Fisher, Matthew P. Webster & Rafael H. Bordini (2012): *Model checking agent programming languages*. *Automated Software Engineering* 19(1), pp. 5–63, doi:10.1007/s10515-011-0088-x. Available at <https://dx.doi.org/10.1007/s10515-011-0088-x>.
- [12] Maria Fox & Derek Long (2003): *PDDL2. 1: An extension to PDDL for expressing temporal planning domains*. *Journal of artificial intelligence research* 20, pp. 61–124, doi:10.1613/jair.1129.
- [13] Javier García, Fern & o Fernández (2015): *A Comprehensive Survey on Safe Reinforcement Learning*. *Journal of Machine Learning Research* 16(42), pp. 1437–1480.
- [14] Michael P Georgeff & Amy L Lansky (1987): *Reactive reasoning and planning*. In: *AAAI*, 87, pp. 677–682.
- [15] Malik Ghallab, Dana Nau & Paolo Traverso (2016): *Automated planning and acting*. Cambridge University Press.
- [16] James Harland, David N Morley, John Thangarajah & Neil Yorke-Smith (2014): *An operational semantics for the goal life-cycle in BDI agents*. *Autonomous agents and multi-agent systems* 28(4), pp. 682–719, doi:10.1007/s10458-013-9238-9.
- [17] James Harland, David N Morley, John Thangarajah & Neil Yorke-Smith (2017): *Aborting, suspending, and resuming goals and plans in BDI agents*. *Autonomous Agents and Multi-Agent Systems* 31(2), pp. 288–331, doi:10.1007/s10458-015-9322-4. Available at <https://dx.doi.org/10.1007/s10458-015-9322-4>.
- [18] Patrik Haslum, Nir Lipovetzky, Daniele Magazzeni & Christian Muise (2019): *An introduction to the planning domain definition language*. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 13(2), pp. 1–187, doi:10.1002/9781118627372.ch2.

- [19] Koen V. Hindriks, Frank S. de Boer, Wiebe van der Hoek & John-Jules Ch. Meyer (2000): *Agent Programming with Declarative Goals*. In: *Proceedings of the 7th International Workshop on Agent Theories and Architectures*, Springer, pp. 228–243. Available at http://dx.doi.org/10.1007/3-540-44631-1_16.
- [20] Xiaowei Huang, Marta Kwiatkowska, Sen Wang & Min Wu (2017): *Safety verification of deep neural networks*. In: *International Conference on Computer Aided Verification*, Springer, pp. 3–29, doi:10.1007/978-3-319-63387-9_1. Available at https://dx.doi.org/10.1007/978-3-319-63387-9_1.
- [21] Viviana Mascardi, Daniela Demergasso & Davide Ancona (2005): *Languages for Programming BDI-style Agents: an Overview*. In: *WOA*, 2005, pp. 9–15.
- [22] D. McDermott, M. Ghallab, A. Howe, C. Knoblock, A. Ram, M. Veloso, D. Weld & D. Wilkins (1998): *PDDL - The Planning Domain Definition Language*. Technical Report TR-98-003, Yale Center for Computational Vision and Control.
- [23] Jonathan Muga & Benjamin Kuipers (2011): *Autonomous learning of high-level states and actions in continuous environments*. *IEEE Transactions on Autonomous Mental Development* 4(1), pp. 70–86, doi:10.1109/TAMD.2011.2160943.
- [24] Hanna M Pasula, Luke S Zettlemoyer & Leslie Pack Kaelbling (2007): *Learning symbolic models of stochastic domains*. *Journal of Artificial Intelligence Research* 29, pp. 309–352, doi:10.1613/jair.2113.
- [25] Anand S Rao & Michael P Georgeff (1992): *An abstract architecture for rational agents*. *KR* 92, pp. 439–449.
- [26] Anand S Rao, Michael P Georgeff et al. (1995): *BDI agents: from theory to practice*. In: *ICMAS*, 95, pp. 312–319.
- [27] Sebastian Sardina & Lin Padgham (2011): *A BDI agent programming language with failure handling, declarative goals, and planning*. *Autonomous Agents and Multi-Agent Systems* 23(1), pp. 18–70, doi:10.1007/s10458-010-9130-9.
- [28] Maarten Sierhuis, William J Clancey & Ron JJ Van Hoof (2007): *Brahms: a multi-agent modelling environment for simulating work processes and practices*. *International Journal of Simulation and Process Modelling* 3(3), pp. 134–152, doi:10.1504/IJSPM.2007.015238. Available at <https://dx.doi.org/10.1504/IJSPM.2007.015238>.
- [29] Richard Stocker, Maarten Sierhuis, Louise Dennis, Clare Dixon & Michael Fisher (2011): *A formal semantics for brahms*. In: *International Workshop on Computational Logic in Multi-Agent Systems*, Springer, pp. 259–274, doi:10.1007/978-3-642-22359-4_18.
- [30] Nicolas Troquard & Laure Vieu (2006): *Towards a Logic of Agency and Actions with Duration*. *Frontiers in Artificial Intelligence and Applications* 141, p. 775.
- [31] Perry Van Wesel & Alwyn E Goodloe (2017): *Challenges in the verification of reinforcement learning algorithms*. Technical Report NASA/TM-2017-219628, NASA Langley Research Center.
- [32] Willem Visser, Klaus Havelund, Guillaume Brat, SeungJoon Park & Flavio Lerda (2003): *Model Checking Programs*. *Automated Software Engineering* 10(2), pp. 203–232, doi:10.1023/A:1022920129859.
- [33] Daniel S Weld et al. (2008): *Planning with durative actions in stochastic domains*. *Journal of Artificial Intelligence Research* 31, pp. 33–82, doi:10.1613/jair.2269.
- [34] Håkan LS Younes & Reid G Simmons (2004): *Solving generalized semi-Markov decision processes using continuous phase-type distributions*. In: *AAAI*, 4, p. 742.