Proceedings of the Twenty-Eighth International Florida Artificial Intelligence Research Society Conference

ORPHEUS: Mental Simulation as Support for Decision-Making in a Virtual Agent

Mihai Polceanu and **Marc Parenthoën** and **Cédric Buche**
Lab-STICC, ENIB
Technopôle Brest-Iroise, 29238 Brest Cedex 3, France.
polceanu@enib.fr, parenthoen@enib.fr, buche@enib.fr

Abstract

Using mental simulation as a means of selecting actions for an agent is not new to the scientific community, but implementations from literature focus on specific scenarios and strategies built by domain-experts. In this paper, we propose a generic decision-making agent architecture which uses mental simulation. Our architecture allows an agent to predict both physical phenomena and behavior of other entities simultaneously, in real time, and to pursue its goal without additional built-in strategies. The experimental results show its performance in a dynamic nature-inspired scenario where our agent changes its behavior based on how it anticipates others will act.

Introduction

Interactive virtual environments pose a wide variety of challenges for intelligent agents, especially to make decisions in order to reach their goals. The difficulty of decision-making tasks rises quickly by introducing continuous space and real time into question (Doya 2000). It also becomes increasingly harder to build intelligent agents that can meaningfully interpret and act in unknown situations.

Classical approaches to decision making, such as production systems (Anderson 1993; Laird 2012), semantic networks and other formal knowledge representation frameworks (Negnevitsky 2005), require domain experts to provide descriptions of the agents' environments, the types of objects to be used and entities to interact with (Castillo 1997). Upon these approaches, learning algorithms have been used to learn new rules (Fürnkranz 1999; Kavšek and Lavrač 2006) and policies (Sutton and Barto 1998) that tell the agent how to behave in various situations. However, their application is limited when dynamic environments are considered (Brooks 1990), where agents must assess multiple interactions between entities and their environment, such as the effects of collision, object shape, action timing and visual occlusion on behavior. Research efforts have been made to address the issues posed by dynamic environments and have yielded important results, such as in robotics (Visser and Burkhard 2007), but challenges still remain that span over several research fields. Concurrently, cognitive architectures have been proposed to integrate multiple techniques into autonomous agents that are better suited for dynamic environments, but some significant aspects such as anticipation and adaptation are still weakly integrated (Vernon, Metta, and Sandini 2007) in existing approaches.

Our contribution is an agent architecture (ORPHEUS[1]), that allows the simulation of a functional model of the world ahead of time, analogous to imagining the outer world and the outcomes of the agent's actions based on the state of the real environment. This paradigm is known to cognitive science as "mental simulation", while the Game AI community uses the term "playout". The novelty of our approach consists in a generic approach to how mental simulations are constructed, namely the ability to simulate physical (such as movement, collision, occlusion) and behavioral (such as mental states, beliefs, decisions) processes within the same framework, as an "imaginary world" (or "sandbox"). The structure of this sandbox world dictates its capability to simulate the evolution of complex interactions between entities and their environment, task at which classic approaches become inapplicable, while modern ones use specific simulators for the task at hand. Our contribution is not a replacement for existing decision making techniques, but a framework in which they can be integrated in an unified fashion with the aim to predict and evaluate possible courses of action, on which decision making can be based.

In this paper, we first present the state of the art related to the use of mental simulation for decision-making, from the perspectives of both cognitive and computer sciences, and motivate the relevance of our contribution. Thereafter, we introduce our proposed generic agent architecture, which is based on the mental simulation paradigm. We then instantiate our architecture into an autonomous agent that is given a goal to achieve in a real time continuous environment, without providing it with information/semantics of its role nor the way in which to behave. We show and discuss aspects of the agent's behavior as a result of using mental simulation, the limitations of the current application of our agent architecture and how these limitations can be overcome.

[1]ORPHEUS: Reasoning and Prediction with Heterogeneous rEpresentations Using Simulation; Source code: https://bitbucket.org/polceanum/orpheus .

Related Work

The mental simulation paradigm enjoys significant interest from the cognitive science community (Kahneman and Tversky 1981; Berthoz 1997; Grezes and Decety 2001). It is used to explain how humans make certain predictions for making decisions, imagining "what if" scenarios (multiple worlds) and revisiting past events in novel ways (Moulton and Kosslyn 2009). Moreover, there exists evidence that mental simulation is not strictly a human capability, but that some animals may also be able to perform it for goal-oriented decision making (Chersi, Donnarumma, and Pezzulo 2013). The principle of mental simulation consists in constructing an imaginary world that can function on its own, based on which various future states of the environment can be inferred and decided upon, resulting in an individual's behavior.

Computational applications of mental simulation are relatively recent and limited to specific scenarios. Results have been obtained by complementing existing systems with prediction and viewpoint adoption capabilities in contexts such as navigation (Bongard, Zykov, and Lipson 2006; Kennedy et al. 2009; Svensson, Morse, and Ziemke 2009; Buche et al. 2010), sensory integration (Cassimatis et al. 2004), object manipulation (Roy, Hsiao, and Mavridis 2004; Kunze et al. 2011; Buche and De Loor 2013), human-agent interaction (Buchsbaum et al. 2005; Breazeal, Gray, and Berlin 2009) and goal recognition (Rao, Shon, and Meltzoff 2004; Gray and Breazeal 2005), which indicate that mental simulation is advantageous over traditional techniques in these chosen scenarios. This is also true for approaches that use heuristic search such as Monte Carlo Tree Search (MCTS) (Uriarte and Ontañón 2014) and Alpha-Beta search (Churchill, Saffidine, and Buro 2012) coupled with simulators designed for the problem at hand. However, in their current state, existing approaches are not applicable to other contexts without considerable adaptation, leaving them highly specific.

Although studies from cognitive science suggest that mental simulation is central to decision making and arguably other important aspects of reasoning, existing approaches do not offer a generic computational model of this paradigm. Therefore, orthogonal to improving existing systems with a specific application of mental simulation, our work focuses on a generic approach to mental simulation which can be in turn applied to given contexts without architectural changes.

Proposed Agent Architecture: ORPHEUS

Our contribution is an agent architecture that provides a generic approach to the paradigm of mental simulation, as support for decision-making. According to this paradigm, an individual owns an imaginary space which is built by sensing reality and functions on its own to predict what outcomes are to be expected in reality in various situations. From the computational standpoint, we consider the "real world" as the environment from which our system takes its input and upon which it acts. In the general sense, the "real world" could be the physical reality. In this work, the "real world" is a virtual environment in which entities can behave autonomously and perceive each other. In addition, we refer to an "imaginary world" as a completely separate virtual space, which is proprietary to the agent(s) using the architecture, where mental simulations and system decision making are performed. In the following, we focus on the structure and functionality of this "imaginary world".

In order for the "imaginary world" to exist, it must be constructed with sensory information. In our approach, perception consists in extracting the state of the "real world" in the form of a *mental image*. The proposed agent architecture bases its decision-making functionality on a continuous cycle of perception, mental simulation and selection process (Fig. 1).

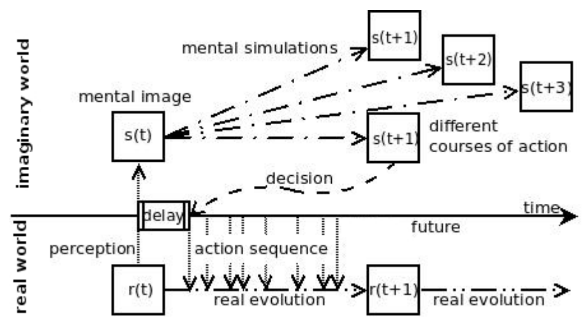

Figure 1: Mental simulation based decision-making.

A mental image statically describes the state of the "imaginary world" at a time t. It is a list of entities (such as objects, birds, etc.), each described by a set of properties (such as position, rotation, shape, linear and angular velocity). The application presented in the following section uses the previously mentioned set of properties due to their direct observability, but within the architecture this set could be modified. We note that non-observable properties of entities can also be represented in this way.

Central to our approach is the process of merging effects from different sources (physics, behavior) into the evolution of the "imaginary world" as a whole. In order to combine the effects of multiple entities on the environment, each entity is assigned models, from an available pool, that dictate the way it behaves. The model assignment for each entity is based on the estimated error between mental simulation and reality, using an error model of choice. Within our architecture, a model can be viewed as a function that takes a state of the world (the assigned entity and its perceptions) as input, and computes a time-dependent change in this state. In this sense, the "imaginary world" falls into the category of multi-agent systems. The process of merging consists in applying a pool of models to their assigned entities in a mental image and accumulating their effects to obtain the subsequent mental image.

A key aspect in the effect merging process in our approach is its granularity. Traditionally, abstract rules are used to describe the evolution of an environment, such as for example: "If [Pencil is not held] Then [new location of Pencil is on Floor]". This type of environment evolution description does not include what happens to the pencil while it falls (obstacles in trajectory), why it falls (gravity), nor what the effects of such process would be, unless specified through a complex array of other rules. In contrast, our approach is to simulate the falling pencil, and therefore naturally represent its interactions with other entities such as the table that may be between the pencil and the floor, or the cat lying on the floor which will run away after the pencil nudges it. In consequence, it is important for this simulation to occur in small time steps, as it is not feasible to combine large effects in the context of complex interactions between entities and their environment. In addition, operations that are not commutative and require an order, such as rotations, can be forced to behave in a nearly commutative way by using small time steps.

We integrate granularity into mental simulations by constructing them from a series of successive steps (Fig. 2). This way, the mental simulation is divided into n steps where, at the i^{th} step, each model (M^j where j is the model identifier) takes the i^{th} mental image as input and computes a small change (δ^j) that is accumulated with the rest of the effects to obtain the $(i+1)^{th}$ mental image.

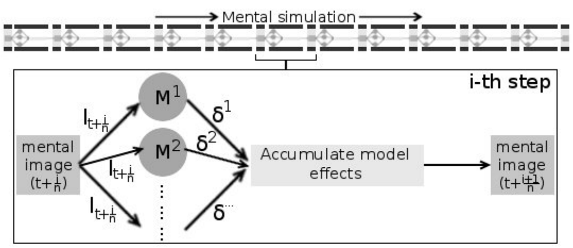

Figure 2: The mechanism through which a mental image is evolved in one step ($i \in [0, n-1]$), by interweaving effects from models.

This approach to constructing mental simulations also allows their evolution to be controlled by integrating different courses of actions for entities and for the agent's mental self. Models that are assigned to entities can also take parameters that control how they generate an entity's behavior which we refer to in this work as "action interfaces". As with the environment and its inhabitants, the agent also includes itself as an entity in the "imaginary world" and assigns models to this mental self. This way, the agent is able to imagine multiple scenarios simultaneously, by manipulating the parameters of its self model in alternative mental simulations.

The agent's decision making process is based on creating an array of different courses of actions ahead of time and, given a goal mental image (goal configuration), providing the list of mental simulations that reach that goal state from which one can later be selected and applied in the "real world". Mental simulations and actions that are performed inside them also embed the notion of the corresponding real time, so that synchronization can be done correctly when the agent performs the actions in the real environment.

Application

Following from the limitations of traditional techniques and related works, we validate our proposed architecture by instantiating it in an autonomous agent which is given a goal in an environment that has the following properties: continuous 3D space, real time, physics (collisions, visual occlusions) and variable behavior of other entities.

Based on the previously mentioned environment property requirements, we chose a scenario (Fig. 3, left) inspired from the predatory behavior of the Felinae subfamily, which consists mostly in small to medium sized cats. It is common for their prey, such as birds, to have wide fields of sight and to be capable of predator detection and avoidance over a distance.

In the natural world, hunting behaviors of cats can be categorized into "mobile" and "stationary", the former being applied in areas abundant in prey while the latter, which consists in ambushing, when the cat is located in areas of interest (Dennis C. Turner 2000). Cats generally employ a stealthy approach followed by a short rush before striking (Kleiman and Eisenberg 1973). In the following, we discuss the instantiation of the generic architecture (Fig. 1) within the context of the chosen test scenario (Fig. 3), aiming to obtain adaptive behavior from the agent through mental simulation, without hard-coding any strategies of its real counterpart; i.e. the agent has no information/semantics that it represents a cat.

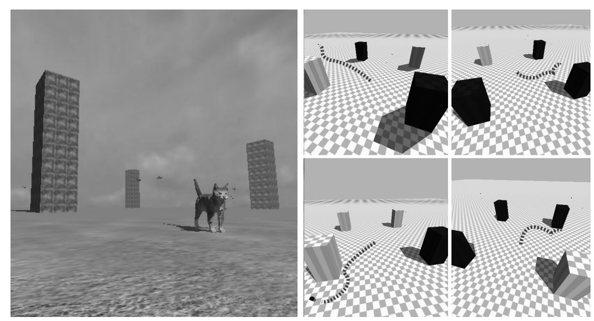

Figure 3: Instantiation of the "real world" (3D environment) featuring the cat and autonomous birds (left), and the "imaginary world" of the cat with examples of mental simulations where it evaluates different action sequences (right).

Experiment Setup

In the "real world" (i.e. the virtual environment that implements the chosen test scenario), birds have boid-like behavior (Reynolds 1987) and are also able to see the cat at an arbitrary distance, given that their vision is not obstructed by an obstacle (birds cannot see through obstacles). A bird that can see the cat will flee (turn away), making it more difficult for the cat to catch it. The experiment consists in six test cases that are constructed by varying obstacle shape (with or without corners) and bird view distance (0 $units$, 50 $units$ and 100 $units$). In the extreme case (100 $units$) the cat is not able to catch birds if no obstacles are present in the environment.

The previously described agent architecture is instantiated in the cat character. The architecture and resulting agent were implemented in C++ with multi-thread support and the tests in this work were run on a Intel® Xeon® 2.80GHz, Ubuntu 12.04 machine with 8GB RAM. We have also tested the system on less powerful machines, where expectedly less simultaneous mental simulations were possible.

We configured the agent to observe the environment (i.e. extract a mental image) at an arbitrary interval of 0.5 seconds to account for real time perception. The mental image is then used to branch an arbitrary number of 20 parallel mental simulations (Fig. 3, right) of varying time lengths (1 to 15 simulated seconds) so as to predict the outcomes of different courses of action into the future. On the test configuration, we obtained a speed ratio of ∼13x; i.e. 15 seconds of mental simulation take ∼1.15 real seconds to compute. When each mental simulation has finished, it is replaced by a new one based on the currently perceived state of the world. A mental simulation will finish when it reaches its time limit or it leads to goal achievement. If there are no mental simulations that lead to goal achievement, the agent will rest still until a solution is found. The number and duration of mental simulations can be varied to increase/decrease chances of a valid solution and to find a near/far solution, respectively.

The agent using our architecture is supplied with three models to evolve mental simulations (Fig. 4). First, a physics model, implemented using Bullet Physics, that applies collision effects and linear and angular velocities on entities. Second, a model that closely approximates the behavior of "real" birds with any of the three possible view distance values. Third, a cat model with parameters which acts only as a control interface. These parameters are a set of possible actions: run forward (values: yes/no), turn (values: left/no/right) and jump (values: yes/no). Model to entity assignments are managed by the agent, so that the error between the behavior of imaginary birds and reality, computed in this application using the Mean Absolute Deviation (Konno and Yamazaki 1991), is minimized.

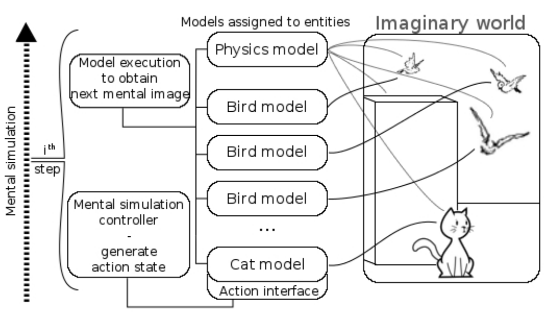

Figure 4: Models assigned to each entity in the imaginary world and action controller for the system's self model (cat). All models contribute to the evolution of the mental simulation at each step.

If at any time the distance between the cat and a bird falls below a threshold value (the cat is required to jump), the bird disappears and is considered caught and the cat's goal is achieved. Once the goal has been satisfied in a mental simulation from the cat's imaginary world, the history of actions performed within it, which correspond to the list of parameter values, are applied in the real world. To avoid any bias from "cleverly chosen" strategies as used by (Buche et al. 2010), the cat's mental simulations include uniformly distributed random parameter values (action sequences). The goal of of this application does not rest in the random action generation itself, as this could be replaced by heuristic search or similar methods, but in illustrating the use of the proposed generic architecture.

Our hypothesis is that through the use of mental simulation, the cat, which is controlled by our system, will adapt its behavior in function of the distance at which birds are able to perceive the danger. By having obstacles in the environment, we hope that the cat will use them to improve its success rate in catching birds, as *this behavior is not explicitly specified*.

Results

Each of the six test cases (Fig. 5) consists of 4 symmetrically placed obstacles and 10 birds which fly in fairly regular circular patterns unless the cat is seen and avoided. Amongst test cases, distances at which the birds can spot the cat are varied and the obstacle shape is also subtly changed.

Results show that the increase in the view distance of birds significantly influences the frequency of locations for the cat, regarding available obstacles. As a general rule, the increase in bird sight determines values in the histograms for the cat (Fig. 5, Subfigures "b") to shift left towards lesser distance from obstacles, and bird catching location frequencies (Fig. 5, Subfigures "c") to invert slope as obstacles prevent birds to see the cat. That is, the agent's behavior consistently adapts to the variation of bird view distance, resulting in a set of distinguishable emerging techniques (Fig. 6).

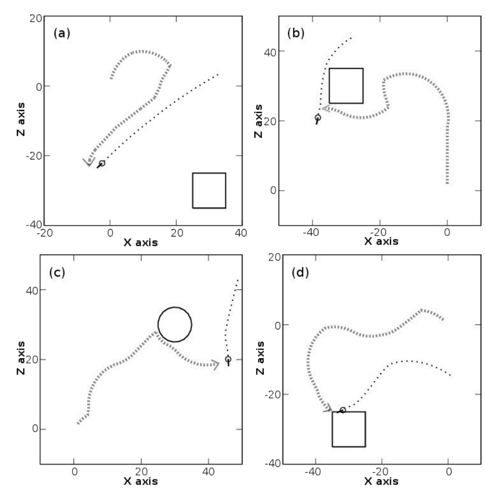

Figure 6: Examples of solutions found by the system for prey catching: (a) simple chase, (b) ambush, (c) hide and chase, (d) blockade. Cat traces shown with gray arrows, and catch locations with directed circles based on bird orientation.

The application of these observed hunting techniques conforms with intuition; namely the cat simply chases (Fig. 6.a) birds with no sight, but when the view distance is increased, it resorts to ambush (Fig. 6.b), hide and chase (Fig. 6.c) and, the strongest example of anticipation observed in the experiment, the blockade (Fig. 6.d) where the agent uses the predicted collision of the bird with an obstacle to catch it.

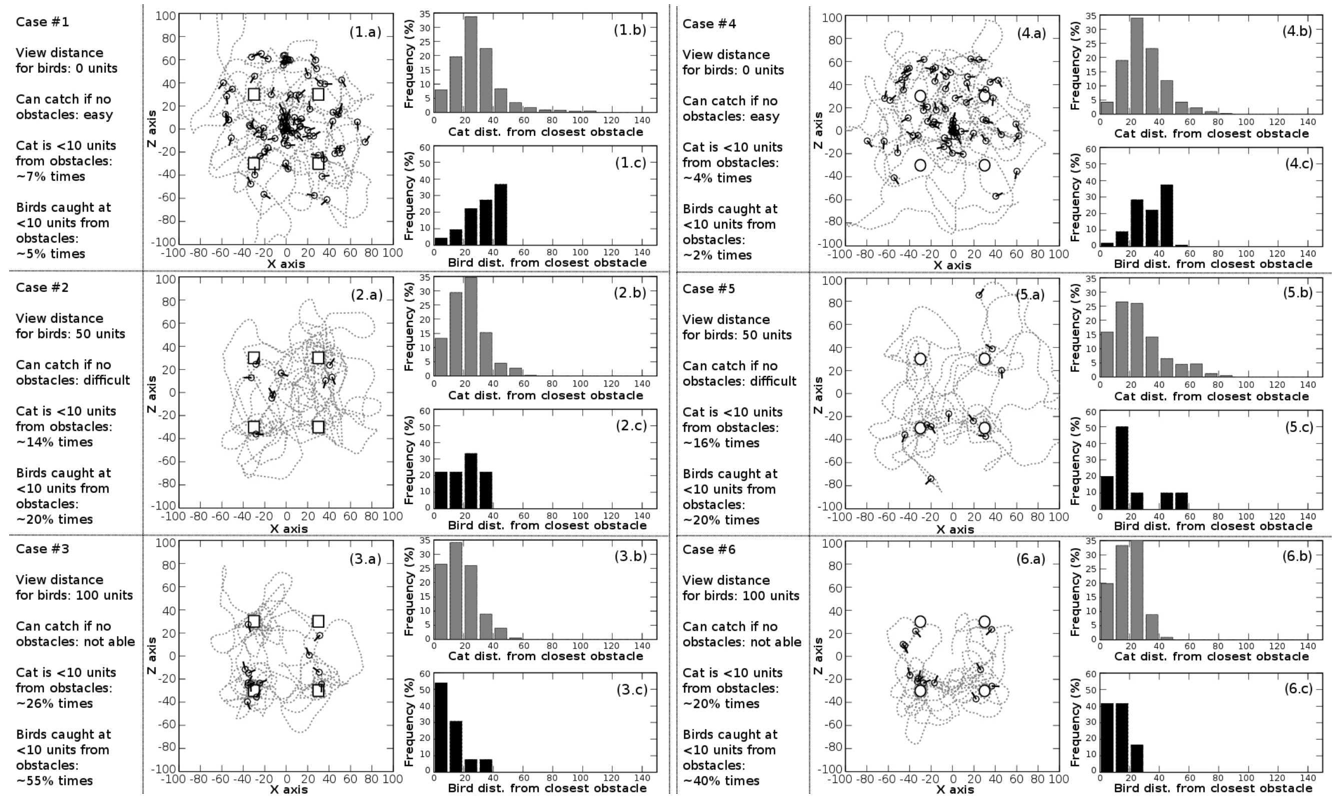

Figure 5: Agent behavior in scenarios with cuboid (cases 1-3) and cylindrical (cases 4-6) obstacles. Subfigures (#.a) illustrate cat (gray) and caught birds (black) traces. Distance histograms are shown for the cat (#.b) and catch locations (#.c) relative to the closest obstacle.

Varying obstacle shape also determines subtle adaptation in our agent, an example of which is illustrated in the 5th test case (Fig. 5, Subfigure 5.c) which features a gap in the cat's position relative to obstacles. This is caused by the increased difficulty of hiding behind cylindrical obstacles which also determines a less prominent use of the ambush technique.

We note that these results are reproducible with a degree of variation in environment configuration such as number of birds and obstacle number and positions, as the agent behaves in the way that leads to goal achievement given different contexts.

Conclusions and Future Work

Our objectives in this work were to propose a generic agent architecture that uses mental simulation as its decision making mechanism, and to verify its functionality within a realistic environment. The ORPHEUS architecture uses a multi-agent imaginary world, in which the instantiated agent can perform mental simulations, to provide a generic interpretation of both physical/environmental and behavioral aspects. This is possible by using multi-step mental simulations that enable the system to interweave the effects from multiple sources, thus merging the two aspects. To validate our approach, we instantiated the proposed generic architecture in an autonomous agent - the cat - and placed it in a virtual environment, with the goal of catching prey. This approach led to adaptiveness to changes in the environment, without predetermined strategies regarding the agent's behavior. The results of the agent within the nature-inspired feline hunting scenario showed that it can exhibit four distinguishable hunting techniques to catch its prey depending on environment conditions, which resemble natural behavior found in felines.

Currently, the application of our system to a wider range of scenarios is limited by the use of inflexible models for prey. In contrast to existing approaches, our architecture supports seamless model replacement, so that it can be applied in other contexts. Special interest is placed on prediction in physical robots. Hence, our research now focuses on replacing these models with time series prediction techniques, without changing the generic architecture that was obtained. Non-linear regression algorithms such as KRLS (Engel, Mannor, and Meir 2003) seem to be well suited for cyclical behavior prediction in the context of mental simulations, and reinforcement learning is a promising approach to improving our agent's efficiency by informing action sequence generation in the imaginary world. These models can be trained simultaneously in real time, creating a pool of candidates that can be selected within our architecture based on their accuracy and used to perform mental simulations, so that the overall performance would benefit from the localized accuracy of heterogeneous sets of learned models.

References

Anderson, J. R. 1993. *Rules of the mind*. Hillsdale, NJ, USA: Lawrence Erlbaum Associates, Inc.

Berthoz, A. 1997. *Le sens du mouvement*. O. Jacob.

Bongard, J.; Zykov, V.; and Lipson, H. 2006. Resilient machines through continuous self-modeling. *Science* 314(5802):1118–1121.

Breazeal, C.; Gray, J.; and Berlin, M. 2009. An Embodied Cognition Approach to Mindreading Skills for Socially Intelligent Robots. *The International Journal of Robotics Research* 28(5):656–680.

Brooks, R. A. 1990. Elephants don't play chess. *Robotics and autonomous systems* 6(1):3–15.

Buche, C., and De Loor, P. 2013. Anticipatory behavior in virtual universe, application to a virtual juggler. *Computer Animation and Virtual Worlds* 24(2):111–125.

Buche, C.; Chevaillier, P.; Nédélec, A.; Parenthoën, M.; and Tisseau, J. 2010. Fuzzy cognitive maps for the simulation of individual adaptive behaviors. *Computer Animation and Virtual Worlds* 21(6):573–587.

Buchsbaum, D.; Blumberg, B.; Breazeal, C.; and Meltzoff, A. N. 2005. A simulation-theory inspired social learning system for interactive characters. In *IEEE RO-MAN 2005*.

Cassimatis, N. L.; Trafton, J. G.; Bugajska, M. D.; and Schultz, A. C. 2004. Integrating cognition, perception and action through mental simulation in robots. *Robotics and Autonomous Systems* 49(12):13–23.

Castillo, E. 1997. *Expert systems and probabilistic network models*. Springer.

Chersi, F.; Donnarumma, F.; and Pezzulo, G. 2013. Mental imagery in the navigation domain: a computational model of sensory-motor simulation mechanisms. *Adaptive Behavior*.

Churchill, D.; Saffidine, A.; and Buro, M. 2012. Fast heuristic search for rts game combat scenarios. In *Proceedings of AIIDE'12*.

Dennis C. Turner, P. B. 2000. IV: Predatory behaviour. In *The Domestic Cat: The Biology of Its Behaviour*. Cambridge University Press. 152–154.

Doya, K. 2000. Reinforcement learning in continuous time and space. *Neural computation* 12(1):219–245.

Engel, Y.; Mannor, S.; and Meir, R. 2003. The kernel recursive least squares algorithm. *IEEE Transactions on Signal Processing* 52:2275–2285.

Fürnkranz, J. 1999. Separate-and-conquer rule learning. *Artificial Intelligence Review* 13(1):3–54.

Gray, J., and Breazeal, C. 2005. Toward Helpful Robot Teammates: A Simulation-Theoretic Approach for Inferring Mental States of Others. In *Proceedings of the AAAI 2005 workshop on Modular Construction of Human-Like Intelligence*.

Grezes, J., and Decety, J. 2001. Functional anatomy of execution, mental simulation, observation, and verb generation of actions: a meta-analysis. *Human brain mapping* 12(1).

Kahneman, D., and Tversky, A. 1981. The simulation heuristic. Technical report, DTIC Document.

Kavšek, B., and Lavrač, N. 2006. APRIORI-SD: Adapting association rule learning to subgroup discovery. *Applied Artificial Intelligence* 20(7):543–583.

Kennedy, W.; Bugajska, M.; Harrison, A.; and Trafton, J. 2009. "Like-Me" Simulation as an Effective and Cognitively Plausible Basis for Social Robotics. *International Journal of Social Robotics* 1(2):181–194.

Kleiman, D., and Eisenberg, J. 1973. Comparisons of canid and felid social systems from an evolutionary perspective. *Animal Behaviour* 21(4):637–659.

Konno, H., and Yamazaki, H. 1991. Mean-absolute deviation portfolio optimization model and its applications to tokyo stock market. *Management science* 37(5):519–531.

Kunze, L.; Dolha, M. E.; Guzman, E.; and Beetz, M. 2011. Simulation-based temporal projection of everyday robot object manipulation. In *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems*, 107–114.

Laird, J. 2012. *The Soar cognitive architecture*. MIT Press.

Moulton, S. T., and Kosslyn, S. M. 2009. Imagining predictions: mental imagery as mental emulation. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences* 364(1521):1273–80.

Negnevitsky, M. 2005. Rule-based expert systems. In *Artificial Intelligence: A Guide to Intelligent Systems*. Pearson Education, 3rd edition. 25–54.

Rao, R. P. N.; Shon, A. P.; and Meltzoff, A. N. 2004. A Bayesian Model of Imitation in Infants and Robots. In *Imitation and Social Learning in Robots, Humans, and Animals: Behavioural, Social and Communicative Dimensions*. Cambridge University Press.

Reynolds, C. W. 1987. Flocks, Herds and Schools: A Distributed Behavioral Model. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '87, 25–34.

Roy, D.; Hsiao, K.-y.; and Mavridis, N. 2004. Mental imagery for a conversational robot. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on* 34(3):1374–1383.

Sutton, R. S., and Barto, A. G. 1998. *Reinforcement learning: An introduction*. MIT press.

Svensson, H.; Morse, A.; and Ziemke, T. 2009. Representation as Internal Simulation : A Minimalistic Robotic Model. In *Proceedings of the Thirty-First Annual Conference of the Cognitive Science Society*, 2890–2895.

Uriarte, A., and Ontañón, S. 2014. Game-tree search over high-level game states in rts games. In *Proceedings of AIIDE'14*.

Vernon, D.; Metta, G.; and Sandini, G. 2007. A survey of artificial cognitive systems: Implications for the autonomous development of mental capabilities in computational agents. *Evolutionary Computation, IEEE Transactions on* 11(2):151–180.

Visser, U., and Burkhard, H.-D. 2007. Robocup: 10 years of achievements and future challenges. *AI Magazine* 28(2):115.