



Simulation within simulation for agent decision-making: Theoretical foundations from cognitive science to operational computer model

Action editor: Peter Erdi

C. Buche^{a,*}, N. Le Bigot^b, M. Polceanu^c

^a ENIB/LAB-STICC, France

^b UBO/LAB-STICC, France

^c FIU/SCIS, USA

Received 12 October 2015; received in revised form 9 February 2016; accepted 9 March 2016
Available online 29 March 2016

Abstract

This article deals with artificial intelligence models inspired from cognitive science. The scope of this paper is the simulation of the decision-making process for virtual entities. The theoretical framework consists of concepts from the use of internal behavioral simulation for human decision-making. Inspired from such cognitive concepts, the contribution consists in a computational framework that enables a virtual entity to possess an autonomous world of simulation within the simulation. It can simulate itself (using its own model of behavior) and simulate its environment (using its representation of other entities). The entity has the ability to anticipate using internal simulations, in complex environments where it would be extremely difficult to use formal proof methods. Comparing the prediction and the original simulation, its predictive models are improved through a learning process. Illustrations of this model are provided through two implementations. First illustration is an example showing a shepherd, his herd and dogs. The dog simulates the sheep's behavior in order to make predictions testing different strategies. Second, an artificial 3D juggler plays in interaction with virtual jugglers, humans and robots. For this application, the juggler predicts the behavior of balls in the air and uses prediction to coordinate its behavior in order to juggle.

© 2016 Elsevier B.V. All rights reserved.

Keywords: Decision-making; Internal simulation; Behavioral simulation; Autonomous behavior; Virtual reality

1. Introduction

For many years, researchers have tended to create virtual environments that provide the opportunity for a human to evolve while interacting with virtual entities. For these virtual worlds to be believable, each entity controlled by the computer must exhibit a behavior giving the illusion of being controlled by another human. This

raises the following question: how can an entity be equipped with believable autonomous behavior in a complex virtual environment in which humans participate?

Traditional symbolic artificial intelligence techniques have been applied to define these behaviors. However, these techniques have limitations as they are mainly based on predetermined rules of behavior chosen by the designer. Indeed, in complex (open simulation, heterogeneous and participatory) virtual worlds, entities may have unpredictable behavior (behavioral variability of autonomous entities, free will of human users), thus creating new situations. When faced with situations unforeseen by the programmer, entities may display unsuitable behaviors.

* Corresponding author.

E-mail addresses: buche@enib.fr (C. Buche), nathalie.lebigot@univ-brest.fr (N. Le Bigot), mpolcean@cs.fiu.edu (M. Polceanu).

In this paper, virtual entities are considered at the same level as human, by integrating human characteristics that are currently lacking in existing artificial intelligence mechanisms. To address this issue, we were inspired by ideas and concepts from cognitive science in defining human decision-making. More precisely, our focus is on finding a computational approach to reproduce adaptive behavior in an intelligent agent, inspired from the human's anticipation ability and capacity to adapt while interacting. We will then be able to examine the use of such concepts into an artificial entity's decision-making process.

This paper is organized as follows. After presenting the concepts from cognitive science which highlight simulation as being an essential aspect of cognition and from which the presented approach is inspired, Section 2 analyzes related works in computational models. 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.

To address this issue, in this paper we present a generic computational model of mental simulation. Thus Section 3 describes a conceptual framework where the entity possesses an autonomous world of simulation within the simulation. In this internal world, the entity is able to simulate itself (using its own model of behavior) and also simulate its environment (using its representation of other entities). The entity also has the ability to anticipate and to learn using internal simulations. In our previous works, the concept of internal simulation has been studied, for test purposes, in two applications described in previous publications (Buche, Chevaillier, Nédélec, Parenthoën, & Tisseau, 2010; Buche & De Loor, 2013; Buche, Jeannin-Girardon, & De Loor, 2011). These applications illustrated the applicability of mental simulation paradigm to decision-making, but connections between decision/anticipation/learning were *ad hoc* for each specific application and required complete architectural modifications to be applied to a new domain.

To test the genericity of the computational model, we have reused these examples in Section 4. The key idea is to show that it is possible, using our architecture, to switch from one domain to the other without making changes to processes binding decision/anticipation/learning. First, we illustrate our approach through an example that simulates dogs gathering sheep. To simulate sheep behavior, the dog uses fuzzy cognitive maps (FCM) of prey. The dog can therefore simulate the sheep's motion in order to make predictions and to test different strategies. Without changing the architecture, we illustrate our proposal by an artificial 3D juggler playing with virtual jugglers, humans and robots. For this application, the juggler predicts the behavior of balls in the air and coordinates its own behavior accordingly, in order to juggle. The virtual juggler uses neural networks to simulate ball motion. The proposed architecture allows the agent to adapt to changes introduced by adding other agents and human users to launch

balls that the virtual juggler can catch while juggling, which was not the case in the original application described in Buche and De Loor (2013). Without changing the architecture, we were able to switch from FCMs to neural networks as a controller for the prediction process. Finally, Section 5 concludes and introduces future work.

2. Context

2.1. Cognitive science toward artificial intelligence

There is a growing body of literature in cognitive science advocating the simulation process as being central in cognition (Decety & Grèzes, 2006; Hesslow, 2002, 2012; Pezzulo, Candidi, Dindo, & Barca, 2013). Contrary to the classical cognitive approach, in these simulative theories, perceptual, cognitive and motor process are not considered as being part of separate domains but rather that sensorimotor processes are fundamental to cognitive activities.

Despite different views of the concept of simulation, one central common point is that simulation corresponds to the reactivation of actions that were formerly executed and stored in memory (Decety & Grèzes, 2006). For example, in his simulation theory (ST), Hesslow (2012) proposed three main assumptions: (1) simulation of action, (2) simulation of perception and (3) anticipation.

The simulation of action (1) implies that brain areas recruited when performing an action are similar to the ones activated when covertly (i.e., when the action is not executed) reactivating the action. In this way, one can consider that a simulated action corresponds to an unexecuted action. The second statement (2) means that perceptual activity may occur in absence of external stimuli. Like during the simulation of action, activation in the brain is similar when perceiving external information and when imaging perceiving this information. Finally (3), Hesslow (2012) proposed the existence of some associative mechanisms allowing both behavioral and perceptual activity that could produce activation in the sensory areas of the brain. The direct consequence is the possibility to generate perceptual activation from simulated action similarly to obtaining this activation from the actual execution of the real action.

In sum, one can simulate both action and perception. When doing so, the recruited brain areas are the same as the ones activated when actually performing the action or actually perceiving external information. Moreover, action simulation can elicit perceptual activity similar to the one which would have occurred if the action were actually performed. The benefit of this anticipation mechanism is two-fold. First, one can be prepared to respond to the consequence of one's own action. Second, one can evaluate in advance the consequence of an action and thus select the most appropriate behavior to achieve the indented goal (Hesslow, 2002; Pezzulo et al., 2013).

The simulation theories are not limited to studying how humans deal with their own actions but they extend to the social domain by taking into account how individuals process other's behavior. The proposition here is that, when observing others acting, one would simulate oneself performing the action (Berthoz, 1997; Jeannerod, 1994). This idea has been reinforced by the discovery, in monkeys, of mirrors neurons which are sensorimotor neurons that fire both when an action is performed and when the same action is observed to be executed by someone else (Rizzolatti, Fadiga, Gallese, & Fogassi, 1996). Since, several observations have supported the existence of a Mirror Neuron System in humans, with shared neural substrates for action identification and action execution ((Rizzolatti & Craighero, 2004) for a review).

Simulating the actions of others would contribute to understanding and anticipating their consequences (Gallese & Goldman, 1998). It has also been proposed that these simulation mechanisms would be significantly involved in more general and higher level social cognition and possibly underlie the capacity of inferring mental states of others, known as Theory of Mind (Gallese & Goldman, 1998; Goldman, 2002). According to the simulation theory (Goldman, 2002, 2005), in order to attribute mental states to others we have to use our own mental mechanisms. In other words, in order to read others' minds, we have to adopt their perspective.

Cognitive science research has also studied how people use simulation mentally represent mechanical systems Hegarty (2004) and reason about how they evolve. Moreover, it has been suggested that there exists a common network in the brain that is responsible for episodic memory, navigation, theory of mind and prospection abilities in humans Buckner and Carroll (2007).

To sum up this section, several propositions in the cognitive science domain consider simulation as a fundamental process of cognition. Through simulation, the brain can anticipate events in the environment, either if they are a consequence of its own action or another's behavior. Thus, simulation mechanisms allow individuals to anticipate future events, so that they can adjust to their ever changing environment. Following this conception, the focus of

our work is to study how internal simulation can enable a virtual agent to adapt to the dynamics of its environment.

An analysis of state of the art computational approaches to internal simulation is provided in the following section. The conclusions of this study allows us to then sketch a conceptual framework which is arguably able to encompass the results of previous approaches and provide a more generic computational solution for building internal simulation based agents. For clarity and tractability, the test scenarios used in this work are considerably simpler than real world human abilities such as reasoning about other people's beliefs and desires, but still preserve highly dynamic behavior, where we evaluate our agents' ability to adapt.

2.2. Computer science related works

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, & Lipson, 2006; Kennedy, Bugajska, Harrison, & Trafton, 2009; Svensson, Morse, & Ziemke, 2009), sensory integration (Cassimatis, Trafton, Bugajska, & Schultz, 2004), object manipulation (Roy, yuh Hsiao, & Mavridis, 2004; Kunze, Dolha, Guzman, & Beetz, 2011), human-agent interaction (Breazeal, Gray, & Berlin, 2009; Buchsbaum, Blumberg, Breazeal, & Meltzoff, 2005) and goal recognition (Gray & Breazeal, 2005; Rao, Shon, & Meltzoff, 2004), which indicate that mental simulation is advantageous over traditional techniques in these chosen scenarios.

In Polceanu and Buche (2013a), we have examined related works in various contexts of using mental simulation for decision-making. Two main directions that can be distinguished in simulation-based approaches to artificial decision-making, anticipation and learning (cf Table 1):

1. Behavioral internal simulation: predict the consequences of actions on behavior to infer knowledge about others

Table 1
Our analysis of the existing computational models.

Computational models objectives	References
<i>1 – Behavioral internal simulation</i>	
1.1 – Imitation of Kin	
– Intention recognition	Buchsbaum et al. (2005), Gray and Breazeal (2005), and Breazeal et al. (2009)
– Learning novelty	Buchsbaum et al. (2005) and Costa and Botelho (2013)
1.2 – Representation of others	
– Predicting behavior	Ustun and Smith (2008), Laird (2001), and Roy et al. (2004)
– Collaboration	Kennedy et al. (2008), Roy et al. (2004), Kennedy et al. (2009), and Hanratty et al. (2007)
– The self	Kunze et al. (2011), Bongard et al. (2006), Dautenhahn and Nehaniv (2002, chap.), Alissandrakis et al. (2002), Gray and Breazeal (2005), Roy et al. (2004), and Kennedy et al. (2008, 2009)
<i>2 – Environmental internal simulation</i>	
2.1 – Virtual environment	
2.2 – Real environment	Kunze et al. (2011), Polceanu and Buche (2013b), Svensson et al. (2009), and Buche and De Loor (2013)
	Roy et al. (2004), Kober et al. (2012), and Cassimatis et al. (2004)

2. Environmental internal simulation: predict the consequences of actions on the physical environment (for example, in the case of a domino scenario, the use of mental simulation would allow the agent to anticipate different sequences of falling pieces)

Taken together, these existing approaches cover all areas of interest in using mental simulation as an anticipation and decision-making technique, however there exists no implementation that addresses all of them on its own. The majority of implementations have begun in virtual reality (Buche & De Loor, 2013; Kunze et al., 2011; Polceanu & Buche, 2013b; Svensson et al., 2009), but fewer have taken the leap to robotic embodiments due to dependencies on specific information about their environment. Those that do however pass into the real world (Cassimatis et al., 2004; Kober, Glisson, & Mistry, 2012; Roy et al., 2004) are either limited in the actions they can perform or rely heavily on repositories of actions that are abstracted for the use within a higher-level framework.

Focus is placed on solving only a subset of the cognitive functions associated with mental simulation, and this is done using specific models of the task at hand. Hence, not many elements are taken into account into mental simulation, for example anticipating trajectories but not collisions (Breazeal et al., 2009; Cassimatis et al., 2004; Kennedy, Bugajska, Adams, Schultz, & Trafton, 2008; Kennedy et al., 2009; Svensson et al., 2009; Ustun & Smith, 2008), or focusing on only one of the environmental and behavioral aspects of the environment, although they are generally interdependent.

To conclude, current approaches are constrained to function in relatively specific setups, far away from a *generic* perspective. Although approaches exist that function in real time, *online learning* is used only for specific tasks. The mechanism of the simulation paradigm, such as continuous imagination-reality comparison in complex environments and imaginative behavior, are not yet fully exploited.

To overcome these shortcomings, in the next section we present a generic framework to model virtual entities that are able to use internal simulation paradigm to anticipate, learn, take decisions and act in dynamic environments where the agent's actions can be influenced or interrupted by other agents or humans.

3. Proposal

Our proposal is based on the theory of internal simulation. The idea is to use this concept as the mechanism to control a virtual entity with the ability to predict and learn. In this section, we first propose a generic framework to model virtual entities (Section 3.1). Constraints and methods regarding the implementation will be discussed afterwards (Section 3.2).

3.1. Generic framework

Following the reasoning expressed in previous sections, we represent three modes to model virtual actor, as follows:

1. **Reactive mode** evolving in the virtual world (Section 3.1.1);
2. **Prediction mode** evolving in the imaginary world (Section 3.1.2);
3. **Learning mode** evolving in the abstract world (Section 3.1.3).

3.1.1. Virtual world

A virtual world is a classical simulation with autonomous entities evolving in interaction. A virtual actor is an agent with sensors for perceiving, effectors for acting and a behavioral model to perform decision-making.

3.1.2. Imaginary world

While acting within the virtual world, each entity can make predictions in its imaginary world. This process consists in simulating the entity's own behavior (with its own behavioral model), and that of the environment it inhabits (using the representation that it has of the behavior of other entities). This simulation occurs a phase ahead of the original simulation, enabling the entities to make predictions. This imaginary space, unique to each entity, functions in parallel with its activity within the virtual world, asynchronously so as not to block the behavioral animation. This imaginary world is a universe of simulation within the simulation (Fig. 1).




Fig. 1. The simulation within simulation process: the entity (left) anticipates the behavior of the other entities (right). It possesses an imaginary world in which it “imagines” what it going to happen.

3.1.3. Abstract world

Prediction in the imaginary world implies a representation of the external world and of its dynamic. To obtain this representation can be a hard challenge because a world adapted to this approach is open: unpredictable interactions can appear at any time and thus the dynamic properties can be disturbed. So, learning mechanisms are a good way to learn the dynamics of the world. In our proposal, predictions in the imaginary world are improved by observing the virtual world online. The virtual actor will then modify its representations of other entities using a learning mechanism. It must be noted that this observed world can also be populated with other actors, or with human-controlled avatars (Stoffregen, Gorday, Sheng, & Flynn, 1999). Similarly, for the approach to be generic, it is important for the control of the behavioral model to be independent of the learning mechanism, so that the model might be piloted by any decisional mechanism. The development of learning adds an entirely new dimension to our model.

3.1.4. Dynamics

In the examples given within this work, virtual entities evolve in a virtual world (first dimension: the virtual world), simulate the representation of behaviors in an imaginary world (second dimension: the imaginary world), and adapt the representation of behaviors through learning (third dimension: the abstract world). The challenge here is therefore to identify the three dimensions and to understand their interactions (Fig. 2). The three worlds evolve in parallel and correspond to three different levels of abstraction. Nevertheless, they are all related and share information. The virtual world provides the necessary

information to the imaginary world in order to simulate an approximate representation of the virtual world. Furthermore, it provides the abstract world with data for adapting its models. The imaginary world feeds back information, particularly concerning the choice of strategies or predictions.

3.2. Implementation

The implementation we propose is sequenced in three steps: (1) observation (2) prediction and (3) learning. Each step is described in the following.

3.2.1. Steps

Step 1: Observation (Models Choice and Estimation)

- **Models Choice.** In most cases of complex dynamic environments, it is not possible to predict all behaviors using only one prediction model. Therefore the question arises of which model is most fit in a given situation. The most common approaches involve an expert who, based on his or her past experience regarding a given problem, can suggest a most likely to succeed solution. However, this type of approach requires the intervention of the human expert when entirely new scenarios arise. To overcome the need for external intervention in model selection, the agent will require to test the prediction accuracy of several models, and choose the most efficient one. In the context of simulation in the simulation paradigm, this model selection is possible due the constant feedback received by the agent from its environment. Having a

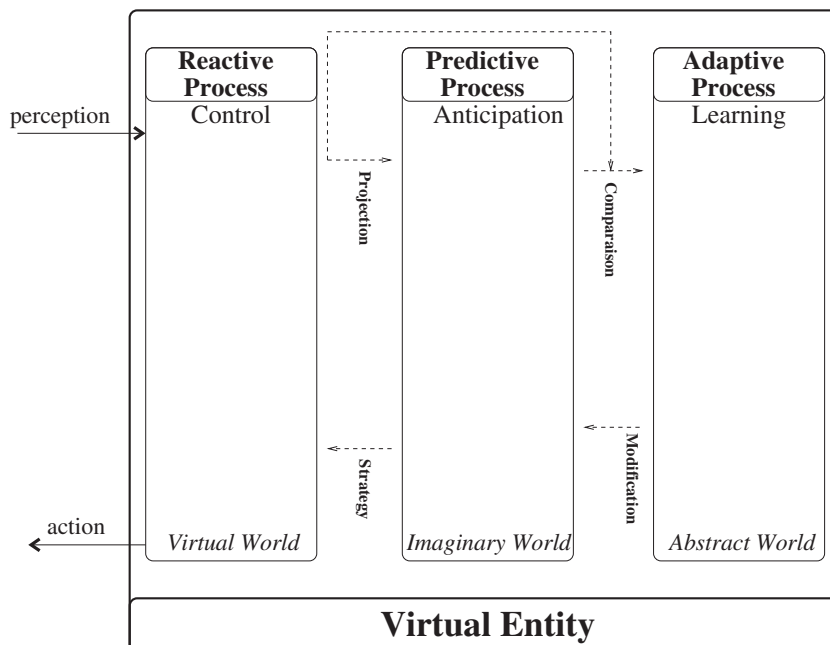


Fig. 2. Representation of the entities' three dimensions: the real world (reactivity mode), the imaginary world (predictability mode) and the abstract world (adaptability mode). The three modes can function asynchronously in parallel.

repository of general behavior models can allow the agent to select different methods for describing each behavior observed in its environment. Thus, we consider the agent has a library of prototypic behaviors models. This library represents the agent’s behavioral culture (Mataric, 2002). For example an animal’s library is made up of the prototypic behavior of both predator and prey.

- **Estimation.** The agent estimates model-sensors and model-effectors through observation. We make the assumption that these features are available.

Step 2: Prediction (Simulating behavior). The agent imagines its own behavior by simulating its own decisional mechanisms and imagines the behavior of the other actors using prototypic behavior models.

Step 3: Learning (Models reconsiderations). The aim is to provide the ability for the agent to adapt its representation of other actors’ behavior. This learning is done using the comparison between the simulation model (imaginary world) and the observation of reality (virtual world).

3.2.2. Technical considerations

Achieving a functional instance of an agent capable of imagining the effects of its actions in the world, and possible outcomes of complex situations, poses a set of technical challenges. The real-time autonomous nature of this agent, raises the question of *execution speed*, as the agent must be able to retrieve the state of the world at any given moment and extract knowledge from this information, to be used within the higher levels of abstraction. In the scenario where the agent interfaces with the real world, the difficulty of this task greatly increases, as computer vision, natural language processing and other interface related aspects still pose a challenge to researchers and developers.

The agent can then create multiple simulations, different versions of reality, through which predictions can be made and solutions to problems can be found. Based on these predictions, the agent can make decisions on how to act, or it can gain novel knowledge about the environment and other inhabitants. Considering that many simulations may be required, a *parallel architecture* is the key approach, so that multiple possibilities can be tested simultaneously, without great performance overhead. Aiming for a network-based architecture which would allow the agent to run simulations on multiple machines would be preferred in this case.

When complexity increases, a complete simulation of the environment might not prove to be efficient, therefore an *attention mechanism* must be put in place to select relevant aspects that require prediction. Moreover, as the simulation happens, events may happen in the real world, and therefore the simulation should be *resynchronized*.

Depending on the speed at which a simulation is run, it can serve different purposes. For instance, a *fast paced simulation* will not be able to receive synchronization data from the real world, because they are no longer at the same virtual time as the simulation. Therefore, to update a fast paced simulation with new data, one must reset it to the current time. These fast paced simulations can be used for short term anticipation, case in which regular reset is acceptable, or in *long term predictions* that do not require input data to be perfectly accurate. Another way to use a simulation is by keeping its virtual time synchronized with the real time. This approach can be used to identify flaws in the prediction models used, or to keep track of some parts of the environment that are not currently observed by the agent.

4. Applications

The framework presented in the previous section is illustrated by two implementations. The first application shows a dog guarding sheep (Section 4.1). The second application shows a juggler playing with a human under several conditions: wind and different types of projectiles (Section 4.2). The scenarios presented in this section are agent-based simulations. The main difference between classical agent-based simulations and our approach consists in the architecture or such agents (i.e. the dog and the juggler) which also perform their own, internal simulation of the simulations they inhabit. Following our proposal, our agents use internal simulation to make their decisions. Simulating the behavior of sheep in an imaginary world allows the dog to test different strategies to gather the sheep. For the juggler, it predicts the trajectory of objects. The abstract world allows to adapt the representation of the behavior of sheep and to adapt the representation of the trajectories of objects. The objective here is not to propose the most effective behavior in a particular context, which could be achieved by conventional agents (Agre & Chapman, 1987), but our aim is to illustrate a generic way to create virtual agents that can function beyond a given context, even if it is seemingly more complicated than that certain context requires.

4.1. Virtual dog gathering sheep

In this application, we implemented a virtual dog gathering virtual sheep (Parentoën, Reignier, & Tisseau, 2001). The motivation is not to simulate an actual dog’s mental processes but to create a virtual agent that can exhibit similar behavior given the context and under dynamic conditions introduced by the human user’s intervention. Starting with a random initial placement of each entity in the experiment (N dog, M sheep and a shepherd), a human user can take control of the shepherd and is able to disturb the simulation.

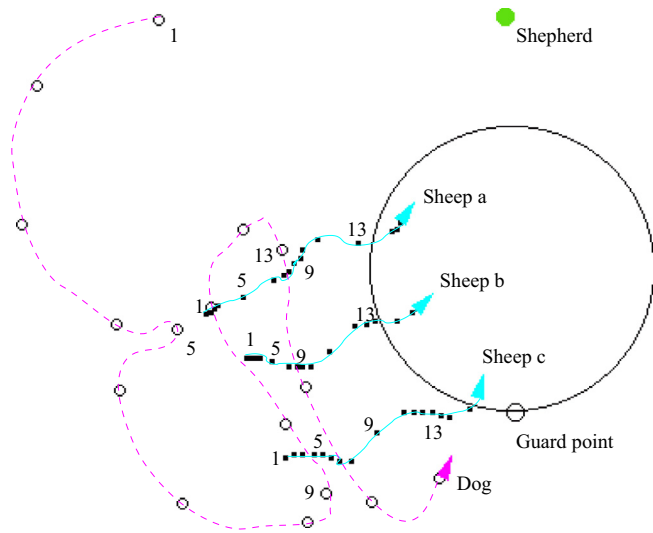


Fig. 3. Trajectory of the sheepdog while bringing back three sheep.

4.1.1. Virtual world

The shepherd moves in the virtual environment. Each sheep distinguishes enemies (dog or human) from friends (sheep) and grass. A sheep has a reserve of energy, increasing while eating and spends while running. By default, it goes straight and ends up becoming exhausted. The sheep eats grass, and becomes afraid of enemies when they are too near. According to the gregarious instinct they socialize. The dog is able to identify the human, sheep, pasture zone and guard point. One or several sheep can move away from the gathering zone. When approaching a sheep, the dog frightens it and obliges it to return to this zone (Fig. 3).

4.1.2. Imaginary world

The dog simulates in its imaginary world several strategies to gather sheep, not by logical reasoning but by a behavioral simulation. Thus, it will be able to make predictions on the future. The dog imagines its behavior by simulating its own decisional mechanism. To simulate sheep

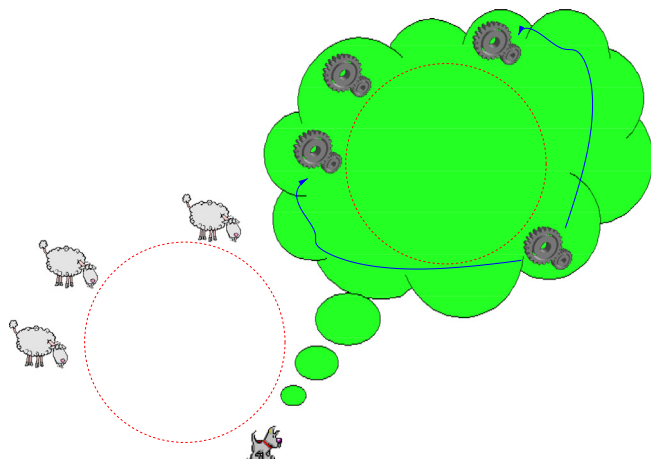


Fig. 4. The dog possesses its own imaginary world in which it can simulate prototypic behavior from a library of behaviors, containing the “prey”.

behavior, the dog uses a model of prey. We choose to use prototypic Fuzzy Cognitive Maps (FCM) (Kosko, 1986) as model of prey (Fig. 4). Modeling predator–prey agents using FCMs has been proposed before by Gras, Devaurs, Wozniak, and Aspinall (2009), however, in our work we focused on learning the FCM weights with the purpose of performing simulation within simulation. Each sheep is associated with its own prototype using FCMs. Thus the dog can simulate sheep behavior and can do predictions. A FCM controls the prototype’s speed and another controls the prototype’s angle (Fig. 5). The dog simulates herding in its imaginary space using two viewing strategies: one associated to vision restricted to a closed neighborhood (dog only takes care about sheep acting in a ten meters radius circle), the other to the largest vision possible (dog takes care about every sheep). The results of these simulations are compared in term of best gathering, then the dog adopts the more suitable strategy to gather sheep. If the herd is divided into two distant groups, the dog afterwards adopts a simulation in its imaginary world the restricted vision strategy to prevent it from running inefficiently between the two groups.

4.1.3. Abstract world

Prototypic FCMs evolve through the learning process described in Buche et al. (2010). The comparison between the result of the imaginary world and the virtual world allows such adaptation in real time. The learning mechanism consists in obtaining the result of the simulation from the imaginary world, comparing it to what happened in the virtual world, and computing an adaptation of prototypic FCMs. We limit our study to the learning of the weights of the causal connections between concepts in a prototypic FCM in order to imitate a given behavior, by modifying neither the structure of the influence graph of a FCM, the fuzzyfication¹ of the sensors, nor the defuzzyfication² of the concepts motors. Fig. 6 illustrates the modification of a prototype, which initially contains default weight values, to adjust the speed of a sheep within the imaginary world to imitate the real sheep’s behavior. We imposed the learning period. Such a period allows the convergence of the process.

By simple observation of the sheep, it estimates information necessary to the fuzzyfication for the prototype. The estimation of sensors values are fuzzyfied in activation of the concepts “Enemy close” and “Enemy far”. The dynamics of the prototype occurs and by defuzzyfication of the activation of the effector motor Escape envy we obtain the image effector. Its corresponds to the representation that the dog has of prey’s speed. This image effector from the prototype is compared to an estimation of sheep’s effectors. This comparison allows to calculate a set of pseudo

¹ Fuzzyfication consists in converting external FCM values to FCM concept activations.

² Defuzzyfication consists in converting FCM concept activation to FCM external values.

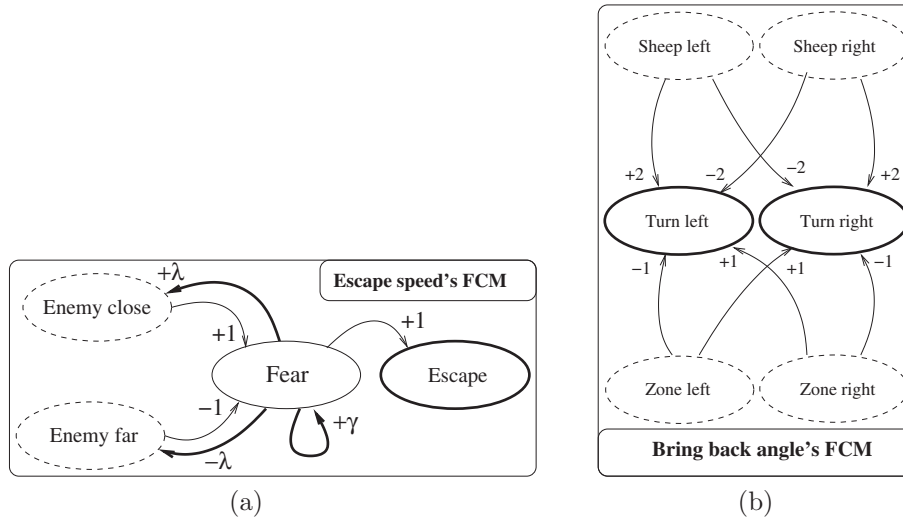


Fig. 5. The prototypic FCMs to represent the sheep controlling speed (a) and angle (b).

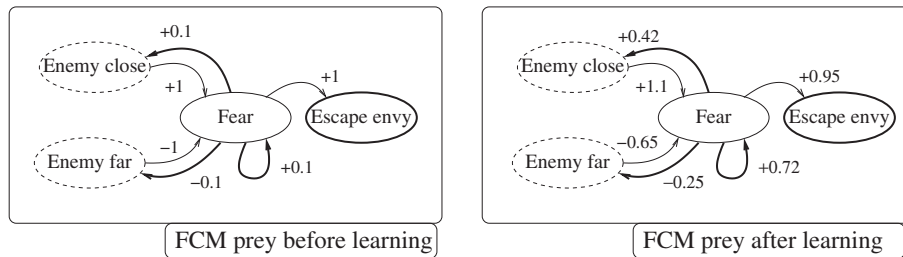


Fig. 6. An FCM of perceptive prey from the library of prototypic FCMs which adapt themselves by learning.

activations that define desired modifications of FCM links. The prey prototype adapts itself to a sheep by reiterating the learning process. Regarding the question of the convergence stability during learning, in practice the experiments undertaken on sheep-dog show that convergence occurs and that the sheep-dog is able to adapt its prototypic FCMs to specific sheep and dogs. We could modify the learning rate through time, as a decreasing sequence tending towards zero. That would ensure a theoretical FCM weight convergence, but the adaptability would be less and less strong with the age of the actor.

4.1.4. Dynamics

Fig. 7 shows the operating principle of an autonomous, predictive and adaptive virtual actor. Three cognitive processes operate in parallel and synchronize only by intermittent messages that lift inhibitions. Roughly, they correspond, from left to right, to short (<1 s), medium (0, 1 s < 10 s) and long term (>10 s) cognition. The first component – the reactive process – operates at high frequency for the acquisition of certain sensor data and is associated with a specified group of cognitive maps for various sensorimotor strategies. The second component – the predictive process – performs internal simulations of average frequency, which can be synchronized with episodic perceptual infor-

mation. This predictive mode can change the reactive one by providing a new sensorimotor strategy. The third component – the adaptive process – compares predictions with perceptions to change behavior patterns used by the other two methods. When prediction error is obvious, the third mode runs in the background at low frequency.

4.1.5. Results

We have implemented three scenarios. First, an adaptation of the dog’s prey prototype to a given sheep, is real time. This application is described in the previous section. Second, the dog learns a way of gathering sheep by the imitation of a human operator or another dog. In this case, the prototypic FCMs used are its own. Third, a “paranoiac” sheep learns how to be surrounded by other sheep and thus remains frightened but does not flee anymore when seeing a dog. Not modifying FCM “paranoiac” links allows to adapt sheep behavior while preserving a “paranoid personality”.

In Fig. 8, we compare the simulation of prototypic sheep behavior in the imaginary world (“prey image”) and the sheep behavior in the virtual world (“sheep model”), before and after learning, while the dog performs the same trajectory. We note that the simulation is closest to reality after learning.

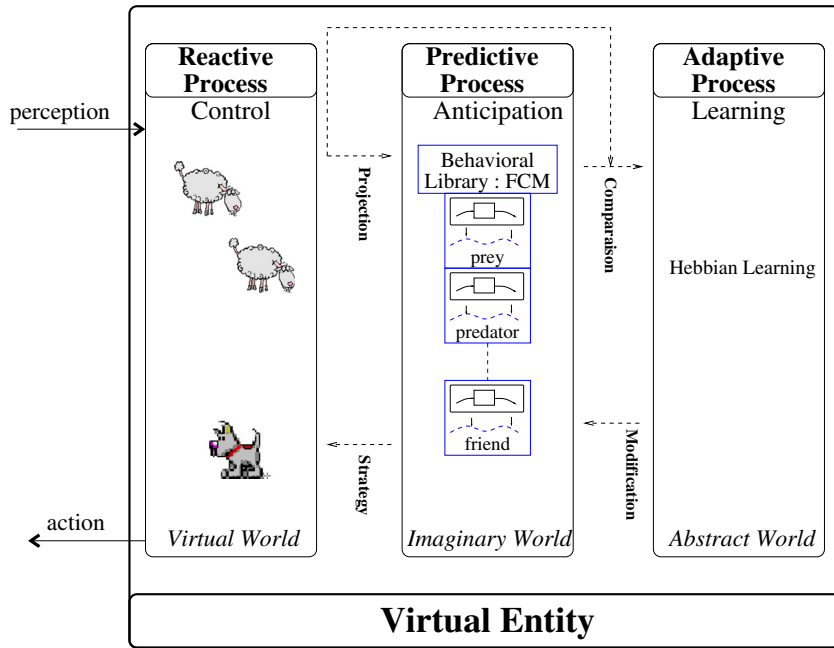


Fig. 7. Instantiation of our framework for a virtual dog.

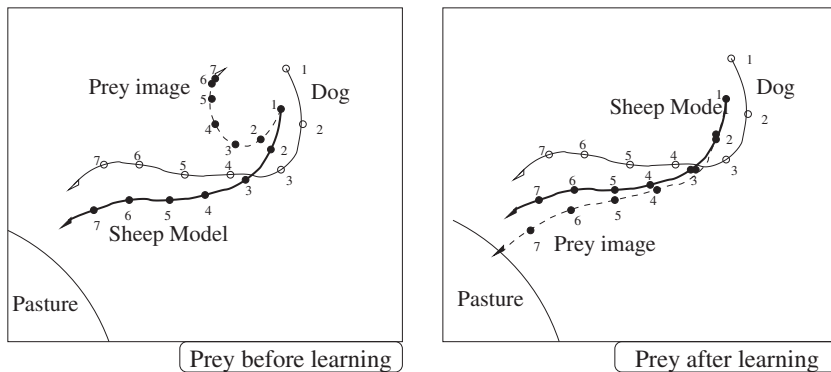


Fig. 8. More pertinent predictions can be obtained from the imaginary world by using imitation learning.

4.2. Virtual juggler

This application shows an artificial juggler which predicts the motion of balls in the air and uses predictions to coordinate its own behavior while juggling. Thanks to this model it is possible to add a human user or a physical robot to launch balls that the virtual juggler can catch while juggling, which was not the case in the original application described in Buche and De Loor (2013).

4.2.1. Virtual world

The virtual world is a universe representing a circus consisting of virtual characters, a human and physical robots juggling together with virtual objects. Robots and humans are physically represented in the virtual world through their avatars (copies of the positions of their bodies).

Concerning virtual juggler, the different phases of juggling are as follows. The juggler begins by looking for a

ball in the air. Once the ball has been spotted, the hand has to be at an estimated reception point (prediction T1). Then, this reception point can be refined. In order to do so, the hand must estimate and correct the anticipated trajectory of the target ball (prediction T2) which is the object of attention. Each hand will therefore be able to catch or miss the target ball. If the ball is caught, the juggler will be able to throw it in the air. Whatever the future of the first ball (caught or missed), the juggler's hand once again starts looking for the next flying ball.

The problem of virtual juggler was discussed in Multon, Ménardais, and Arnaldi (2001) and Julliard and Gibet (1999). But, in these approaches, neither the modeling of approximate anticipation nor the theory of internal simulation was taken into account. More generally, the relationships between cognitive sciences and character's behavior were not addressed.




Fig. 9. Juggler predicts the motion of the balls in the air and uses its predictions (T1 and T2) to coordinate its own behavior in order to continue to juggle.

4.2.2. Imaginary world

While acting in the virtual environment, the juggler predicts the trajectories of projectiles in its imaginary world. The objective is to predict the evolution of their path by approximating the physics and also the perturbations that influence it, such as variable wind speed. The approximate position of the balls (T1 and T2) is obtained through their simulation in the imaginary world of juggling. Within the context of juggling, information must be gathered quickly in order to maintain the juggling dynamics. The use of mathematical functions from perceptron-type neural networks (NNs) to make predictions about the trajectory is adequate. Furthermore, NNs correspond to the need to manipulate (both spatial and temporal) digital data (see Fig. 9).

It is, of course, also possible to use deterministic equation models of movement to make predictions. However, such precise predictions would be extremely noise-sensitive (disruption of the environment as the ball falls) and would not account for the use of approximations and readjustments in real-time which seem to be the basis of the anticipatory mechanisms that we aim to respect (Berthoz, 1997).

4.2.3. Abstract world

The abstract world corresponds to the learning process of weights of the arcs of these networks (Buche et al., 2011). Since they are universal approximators, they allow real-time adaptation of the juggler gestures to different types of disturbances.

We chose a topology with two hidden layers as the aim was to approximate a continuous function (Cybenko, 1989). Each hidden layer has 19 neurons, and we thus obtain $3 \times 19 \times 19 \times 3$ multilayer perceptrons. We assign the perceptron weights with given values prior to learning. The activation function of the neurons is limited. The learning algorithm is a backpropagation of the gradient error. Learning is thus conducted with a maximum of

100 iterations using the Fast Artificial Neural Network³ (FANN) library.

4.2.4. Dynamics

Fig. 10 shows the operating principle of an autonomous, predictive and adaptive virtual actor. Three cognitive processes operate in parallel and synchronize only by intermittent messages that lift inhibitions.

The general features of this proposition allow interactions between several jugglers. To do that, the only change is the direction of the ball launched by each juggler (Fig. 11a). The juggler can also catch a new ball thrown by a human user (Fig. 11b). This is useful for evaluating the believability of the virtual juggler (real-time decision-making, online adaptation, etc.). The human user interacts with the virtual juggler by using different devices (1) wii-mote (2) data gloves (3) razer hydra or (4) arm with force feedback.

The virtual juggler can also juggle with an avatar of a physical robot (Fig. 12). The environment then provides interactions between virtual jugglers, humans and physical robots (Fig. 13). The decision of the robot uses the same mechanisms (simulation within simulation) as the virtual juggler does.

4.2.5. Results

We evaluate the anticipatory mechanism regarding its qualities (impact on decision-making) and the final result (the juggler animation). Regarding the quality, the generalization abilities of NN allow the online adaptation of the juggler's motion to disturbances. Two experiments consist in disturbing the juggler to validate its robustness to variability in the environment. At first, noise is introduced in the projectile trajectories: we exchange balls by maces (Fig. 14a). Through the prediction by NN T1 is less accurate, NN T2 is able to correct it properly, and the juggler continues to juggle when balls are transformed in maces. Second, gravity in the virtual environment is modified, and wind is added (Fig. 14b). The juggler is not informed of these changes. Thanks to the use of NN, our juggler can adapt its arm's position in a precise position, and it can catch a ball even if some small disturbances arrive. Classical planning approaches are used to face these problems: if they use discrete variables, it implies the impossibility to adopt continuous values (such as the arm position); if they use approximate reasoning – which allows us to correct in line a drift from a planning it implies that this drift should be provided and then explicitly represented.

Regarding the final result, the virtual juggler has the abilities to adapt and to learn the changing behavior of the environment online. Of course, this is also the case with classical reinforcement learning algorithms, but with this kind of algorithms, one learns qualities associated to discrete states. This mechanism takes a very long time and

³ <http://leenissen.dk/fann/wp/>.

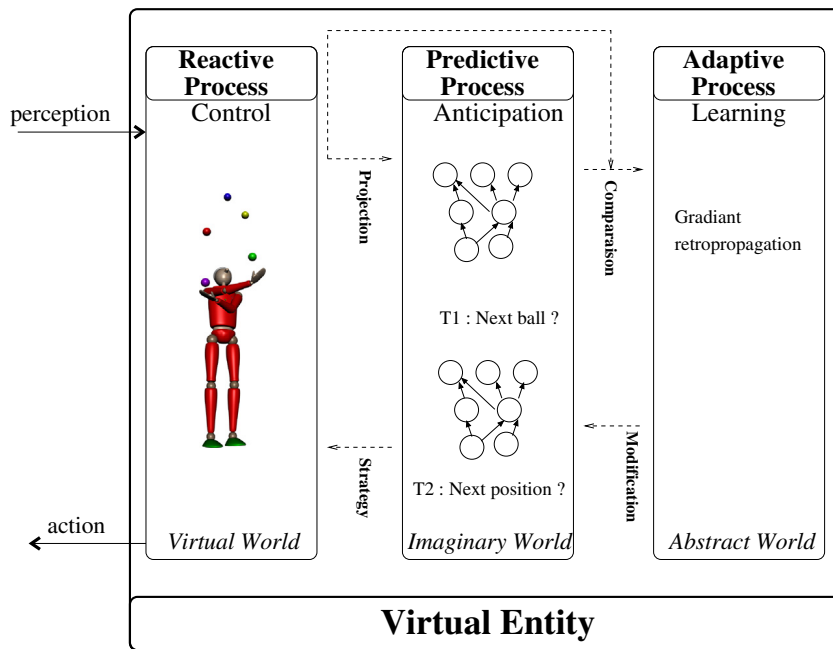


Fig. 10. Instantiation of our framework for a virtual juggler.



(a)



(b)

Fig. 11. Multi-jugglers (a) and a human can juggle with the virtual juggler using the Wiimote (b).



(a)



(b)

Fig. 12. Physical robot (a) juggling with virtual juggler (b).



Fig. 13. Physical robot, human and virtual jugglers juggling.

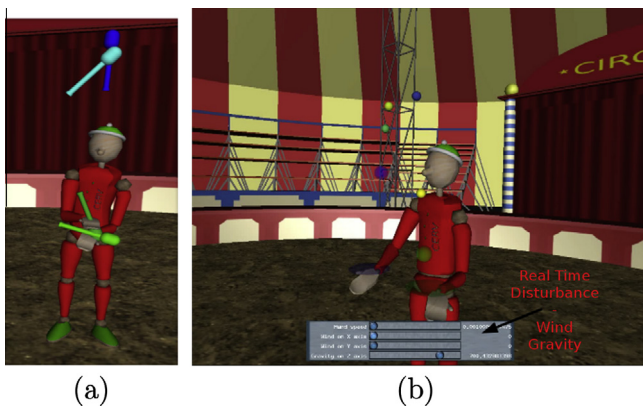


Fig. 14. Juggling with maces (a) and real time disturbance (b).

is not tractable online during an interaction with a human. Moreover, it generally addresses discrete decision when, for our juggler, the decision leads to a precise position of the hand. In our case, the system does not only recognize changes in behavior but can also learn these changes. To sum up, without our architecture, the juggler would be unable to credibly lose its ball and to adapt to a human juggler, which never interacts exactly in the same way because it is the human nature for behavior to be imprecise but rarely irrelevant. These experiments and the final result can be watched at <https://youtu.be/rOZPPRLEUyg>.

5. Conclusion and future works

5.1. Proposal

For the behavioral believability of the interaction of a virtual entity to increase, it would seem essential to integrate an anticipatory capacity by which the behavior of other entities and their consequences on the environment can be predicted. To do so, we suggest an architecture by which the three modes – reactivity, predictability, and

adaptability – can function asynchronously in parallel. The prediction is made by an autonomous world of a simulation within a simulation, in which the entity can simulate itself (with its own behavioral model) and its environment (with the representations that it constructs of the behaviors of other entities).

A first application is related to a real-life example involving different types of agents: a shepherd, dogs, and a herd of sheep. The simulation within the simulation process allows the dog to simulate sheep behavior (using FCMs) and to anticipate the result of different strategies to gather sheep. In the abstract world, the learning mechanism allows the dog to adapt an FCM prey prototype to a given sheep in real time, which leads to their predictions becoming more significant. In a second application, we developed a virtual juggler that anticipates the trajectory of the balls without calculating them precisely. The juggler hypothesizes using an open and uncertain environment with variable properties, that is to say, that are unknown from an analytical standpoint. We therefore use universal approximators obtained through learning.

5.2. Limits

The drawbacks of all these important properties are as follows: (1) the lack of correlation with real data and (2) the need for dedicated models of predictions.

These drawbacks indicate how some work still remains to address a real autonomous agent.

5.3. Outlook

The next steps in applying our approach to more diverse scenarios, consist in further research on overcoming the previously mentioned limitations that it faces. It is clear that models that are fine-tuned for specific scenarios will cease to function properly when applied to completely novel situations. Therefore, it is important for the agent using our approach to have the capability of constructing new models in real time, using a more flexible repository of learning techniques. Furthermore, for the agent to be truly autonomous, the human expert that chooses which models are applied to each scenario has to be removed from the loop, therefore the agent requires a mechanism to choose which are the current best models to apply. This can be achieved in a generic way through a constant “reality check”, by which the internal simulations are evaluated in comparison with the real evolution of the environment and their accuracy can be ranked accordingly. Our research now focuses on developing the generic architecture entitled “ORPHEUS: Reasoning and Prediction with Heterogeneous rRepresentations Using Simulation” (Polceanu, 2015) with the aim to account for the limitations of our current results, that would lead to a more complete use of the “simulation within simulation” paradigm.

References

- Agre, P., & Chapman, D. (1987). Pengi: An implementation of a theory of activity. In *Proceedings of the six national conference on artificial intelligence* (pp. 268–272).
- Alissandrakis, A., Nehaniv, C., & Dautenhahn, K. (2002). Imitation with alice: Learning to imitate corresponding actions across dissimilar embodiments. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, 32, 482–496.
- Berthoz, A. (1997). Le sens du mouvement. O. Jacob.
- Bongard, J., Zykov, V., & Lipson, H. (2006). Resilient machines through continuous self-modeling. *Science*, 314, 1118–1121.
- Breazeal, C., Gray, J., & Berlin, M. (2009). An embodied cognition approach to mindreading skills for socially intelligent robots. *The International Journal of Robotics Research*, 28, 656–680.
- Buche, C., Chevaillier, P., Nédélec, A., Parenthoën, M., & Tisseau, J. (2010). Fuzzy cognitive maps for the simulation of individual adaptive behaviors. *Computer Animation and Virtual Worlds (CAVW)*, 21, 573–587.
- Buche, C., & De Loor, P. (2013). Anticipatory behavior in virtual universe, application to a virtual juggler. *Computer Animation and Virtual Worlds (CAVW)*, 24, 111–125.
- Buche, C., Jeannin-Girardon, A., & De Loor, P. (2011). Simulation theory and anticipation as a basis for interactive virtual character in an uncertain world. application to a human-virtual characters interaction for juggling. *Computer Animation and Virtual Worlds (CAVW), Computer Animation and Social Agents (CASA'11) Special Issue*, 22, 133–139.
- Buchsbaum, D., Blumberg, B., Breazeal, C., & Meltzoff, A. (2005). A simulation-theory inspired social learning system for interactive characters. In *IEEE international workshop on robot and human interactive communication, 2005. ROMAN 2005* (pp. 85–90).
- Buckner, R. L., & Carroll, D. C. (2007). Self-projection and the brain. *Trends in Cognitive Sciences*, 11, 49–57.
- Cassimatis, N. L., Trafton, J. G., Bugajska, M. D., & Schultz, A. C. (2004). Integrating cognition, perception and action through mental simulation in robots. *Robotics and Autonomous Systems*, 49, 13–23.
- Costa, P., & Botelho, L. (2013). Learning by observation of agent software images. *Journal of Artificial Intelligence Research*.
- Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2, 303–314.
- Dautenhahn, K., & Nehaniv, C. L. (2002). *Imitation in animals and artifacts. The agent-based perspective on imitation*. Cambridge, MA, USA: MIT Press, pp. 1–40.
- Decety, J., & Grèzes, J. (2006). The power of simulation: Imagining one's own and other's behavior. *Brain Research*, 1079, 4–14.
- Gallese, V., & Goldman, A. (1998). Mirror neurons and the simulation theory of mind-reading. *Trends in Cognitive Sciences*, 2, 493–501.
- Goldman, A. (2002). Simulation theory and mental concepts. *Simulation and Knowledge of Action*, 119.
- Goldman, A. (2005). Imitation, mind reading, and simulation. *Perspective on Imitation, from Neuroscience to Social Science*, 2, 7993.
- Gras, R., Devaurs, D., Wozniak, A., & Aspinall, A. (2009). An individual-based evolving predator-prey ecosystem simulation using a fuzzy cognitive map as the behavior model. *Artificial Life*, 15, 423–463.
- Gray, J., & 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*.
- Hanratty, T., Hammell, R. J., II, Yen, J., & Fan, X. (2007). Utilizing concept maps to improve human-agent collaboration within a recognition-primed decision model. In *2007 IEEE/WIC/ACM international conference on Intelligent Agent Technology (IAT'07)* (pp. 116–120).
- Hegarty, M. (2004). Mechanical reasoning by mental simulation. *Trends in Cognitive Sciences*, 8, 280–285.
- Hesslow, G. (2002). Conscious thought as simulation of behaviour and perception. *Trends in Cognitive Sciences*, 6, 242–247.
- Hesslow, G. (2012). The current status of the simulation theory of cognition, 1428, 71–79.
- Jeannerod, M. (1994). The representing brain: Neural correlates of motor intention and imagery. *Behavioral and Brain Sciences*, 17, 187–202.
- Julliard, F., & Gibet, S. (1999). Reactiva'motion project: Motion synthesis based on a reactive representation. In *GW '99: Proceedings of the international gesture workshop on gesture-based communication in human-computer interaction* (pp. 265–268). London, UK: Springer-Verlag.
- Kennedy, W. G., Bugajska, M. D., Adams, W., Schultz, A. C., & Trafton, J. G. (2008). Incorporating mental simulation for a more effective robotic teammate. In *Proceedings of the 23rd national conference on artificial intelligence. AAAI'08* (Vol. 3, pp. 1300–1305). AAAI Press.
- Kennedy, W., Bugajska, M., Harrison, A., & Trafton, J. (2009). like-me simulation as an effective and cognitively plausible basis for social robotics. *International Journal of Social Robotics*, 1, 181–194.
- Kober, J., Glisson, M., & Mistry, M. (2012). Playing catch and juggling with a humanoid robot. In *IEEE-RAS international conference on humanoid robots*.
- Kosko, B. (1986). Fuzzy cognitive maps. *International Journal of Man-Machine Studies*, 24, 65–75.
- Kunze, L., Dolha, M. E., Guzman, E., & Beetz, M. (2011). Simulation-based temporal projection of everyday robot object manipulation. In *Proc. of 10th int. conf. on autonomous agents and multiagent systems* (pp. 107–114).
- Laird, J. E. (2001). It knows what you're going to do: Adding anticipation to a quakebot. In *Proceedings of the fifth international conference on autonomous agents* (pp. 385–392).
- Mataric, M. (2002). Visuo-motor primitives as a basis for learning by imitation: Linking perception to action and biology to robotics. In K. Dautenhahn & C. Nehaniv (Eds.), *Imitation in animals and artifacts* (pp. 392–422). MIT Press.
- Multon, F., Ménardais, S., & Arnaldi, B. (2001). Human motion coordination: A juggler as an example. *The Visual Computer*, 17, 91–105.
- Parenthoën, M., Reignier, P., & Tisseau, J. (2001). Put fuzzy cognitive maps to work in virtual worlds. In: *IEEE international conference on fuzzy systems (fuzz-IEEE)* (Vol. 1, pp. 252–255). Melbourne, Australia.
- Pezzulo, G., Candidi, M., Dindo, H., & Barca, L. (2013). Action simulation in the human brain: Twelve questions. *New Ideas in Psychology*, 31, 270–290.
- Polceanu, M. (2015). ORPHEUS: Reasoning and Prediction with Heterogeneous rRepresentations Using Simulation. Ph.D. thesis Université de Bretagne Occidentale (UBO).
- Polceanu, M., & Buche, C. (2013a). Towards a theory-of-mind-inspired generic decision-making framework. CoRR, abs/1405.5048.
- Polceanu, M., & Buche, C. (2013b). Towards a theory-of-mind-inspired generic decision-making framework. In *IJCAI 2013 Symposium on AI in Angry Birds*.
- Rao, R. P. N., Shon, A. P., & 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.
- Rizzolatti, G., & Craighero, L. (2004). The mirror-neuron system. *Annual Review of Neuroscience*.
- Rizzolatti, G., Fadiga, L., Gallese, V., & Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3, 131–141.
- Roy, D., yuh Hsiao, K., & Mavridis, N. (2004). Mental imagery for a conversational robot. *IEEE transactions on systems, man, and cybernetics, part B: cybernetics*, 34, 1374–1383.
- Stoffregen, T., Gorday, K., Sheng, Y.-Y., & Flynn, S. (1999). Perceiving affordances for another person's actions. *Journal of Experimental Psychology: Human Perception and Performance*, 25, 120–136.
- Svensson, H., Morse, A., & Ziemke, T. (2009). Representation as internal simulation: A minimalistic robotic model. In *Proceedings of the thirty-first annual conference of the cognitive science society* (pp. 2890–2895).
- Ustun, V., & Smith, J. (2008). Mental simulation for creating realistic behavior in physical security systems simulation. In *Simulation conference, 2008. WSC 2008. Winter* (pp. 879–885).