



# Robotic Simulations, Simulations of Robots

Edoardo Datteri<sup>1</sup> · Viola Schiaffonati<sup>2</sup>

Received: 22 May 2018 / Accepted: 16 January 2019 / Published online: 25 January 2019  
© Springer Nature B.V. 2019

## Abstract

Simulation studies have been carried out in robotics for a variety of epistemic and practical purposes. Here it is argued that two broad classes of simulation studies can be identified in robotics research. The first one is exemplified by the use of robotic systems to acquire knowledge on living systems in so-called biorobotics, while the second class of studies is more distinctively connected to cases in which artificial systems are used to acquire knowledge about the behaviour of autonomous mobile robots. The two classes pertain to sub-areas of robotics which are apparently quite distant from one another in terms of goals, methodologies, technologies, and theoretical backgrounds. Still both are concerned with building, running, and experimenting on simulations of other systems. This paper aims to reveal and discuss some methodological commonalities between the two classes of studies. Philosophical literature on simulation methodologies has been traditionally focused on studies carried out in research fields other than robotics: this article may therefore contribute to shedding light on how the concept of simulation is used in robotics, and on the role simulation methodologies play in this research field.

**Keywords** Robotic simulation · Simulation system · Biorobotics · Autonomous robotics

## 1 Introduction

The term “simulation” is often adopted, in scientific and philosophical literature, to denote computational devices used to acquire knowledge about physical systems, as exemplified by computer simulations of atmospheric processes. Simulation studies

---

✉ Viola Schiaffonati  
viola.schiaffonati@polimi.it

<sup>1</sup> RobotiCSS Lab, Laboratory of Robotics for the Cognitive and Social Sciences, Department of Human Sciences for Education, Università degli Studi di Milano-Bicocca, Piazza dell’Ateneo Nuovo 1, 20126 Milan, Italy

<sup>2</sup> Artificial Intelligence and Robotics Laboratory, Department of Electronics, Information, and Bioengineering, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milan, Italy

have been often carried out in robotics too. Indeed, here it is argued that two broad classes of simulation studies can be identified in robotics research. The first class of studies is exemplified by robotic systems used to acquire knowledge about living systems, while, in the second class of studies, artificial systems (programmed computers) are used to acquire knowledge about the performance of robotic systems. The two classes pertain to sub-areas of robotics—biorobotics and autonomous mobile robotics—which are apparently distant from one another, with little overlapping in terms of goals, methodologies, technologies, and theoretical backgrounds. Still, they are both concerned with building, running, and experimenting on simulations of other systems: this paper aims to reveal and discuss some methodological commonalities between the two classes.

This goal is worth pursuing for many reasons. First, as pointed out before, philosophical literature on simulations has been traditionally focused on research fields other than robotics, typical case-studies being computer simulations of atmospheric, economic, physical and social phenomena (Weisberg 2013).<sup>1</sup> This article may therefore contribute to shedding light on the epistemic role of simulations in robotics. Second, the conceptual framework developed here may contribute to defining the notion of *robotic* simulation and to outlining the methodological structure of simulation studies carried out in robotics. Third, by revealing procedural commonalities between what is done in relatively distant provinces of robotics, one may pinpoint methodological problems arising in both areas, and extend solutions developed in one area to the other one.

The structure of this article is as follows. In Sect. 2, some representative case-studies will be described concerning biorobotics (Sect. 2.1) and autonomous mobile robotics (Sect. 2.2). Section 3 will provide clarifications on how some key concepts, such as “simulation” and “theoretical model”, are used in this article. This section will also reveal that the case-studies described in Sect. 2 share a common methodological structure, all being cases in which a simulation system is used to acquire knowledge about a target system. In Sect. 4 some methodological peculiarities of each class will be discussed, based on the case-studies presented. It will be pointed out that in autonomous mobile robotics, but not in biorobotics, the target system can be a fictional entity, and that comparisons between the behaviour of the simulation system and the target system are carried out in the two classes of studies, albeit with different purposes.

---

<sup>1</sup> There is a limited methodological literature on simulations in biorobotics (Datteri 2017; Datteri and Tamburrini 2007; Webb 2001, 2006), and no methodological literature on the role of simulations in autonomous mobile robotics.

## 2 Biorobotics and Autonomous Mobile Robotics: Some Case-Studies

### 2.1 Biorobotics: On Lobsters and Plesiosaurs

Biorobotics is a branch of robotics concerned with the use of robotic systems as simulations to acquire knowledge about living organisms, a use that dates back to cybernetics (Rosenblueth and Wiener 1945) and before (Cordeschi 2002). Biorobotics has received renewed interest in the last decades, especially after the proposal and the theoretical systematization of behaviour-based robotics (Arkin 1998) and the subsequent attempts to mimic the behaviour of insects and animals in biologically inspired robotics (Pfeifer 2009). As often pointed out in the literature, the purpose of biorobotics is not merely to draw inspiration from living systems to build more efficient robots: the goal is rather to build robots which may serve as experimental tools to study the behaviour of living systems. A couple of examples will be useful to understand this point.

The first one is relatively classic in the biorobotics literature. The goal of the study described in Grasso et al. (2000), and called “lobster study” from now on, is to explain how lobsters can reach sources of chemical streams, possibly signalling the presence of food, in the water. The authors start with a very simple mechanistic hypothesis. Imagine the chemical stream propagating from the source. It is reasonable to suppose that chemical concentration will be high near the source and low at the periphery, and that, consequently, local differences in chemical concentration will reflect the relative position of the source. If this were the case, the animal would just have to climb the concentration gradient in order to reach the destination. Lobsters’ antennae are chemical sensors. The authors therefore hypothesize that lobsters, to perform the task, steer towards the side corresponding to the higher concentration level based on the following relationship: the higher the concentration perceived at the left antenna is, the higher the velocity of the right motor organs of the animal will be, and vice versa. This simple mechanism is conceptually akin to the structure of one of the first adaptive agents described by Valentino Braitenberg (1986).

This mechanistic hypothesis is intuitively plausible—at least, under the assumption that the chemical stream spreads in a regular way from the source. But water is turbulent, and turbulence may perturb chemical stream at the periphery. If this were the case, local differences in chemical concentration in the periphery would be poorly informative of the position of the source, and the proposed mechanism would probably disorient the animal. Is the chemical stream regular enough to allow the simple mechanism proposed by the authors to guide the animal towards the destination?

*RoboLobster* is a small underwater robot (24 cm in length) moving on wheels and provided with two sensors, one on each side, able to detect chemical concentration (chemoreceptors). The control mechanism reflects the hypothesis above: each lateral chemoreceptor is connected to the contro-lateral motor, speeding it up when higher concentration is detected, as prescribed by the hypothesized mechanism. The idea is to put the robot in the water within a chemical stream

propagating from a distant source: if the robot consistently succeeds in reaching the chemical source (thus matching lobsters' chemiotaxis abilities), then one may be induced to corroborate the mechanistic hypothesis under scrutiny. Otherwise, one may be induced to reject it. The authors carry out these experiments and obtain non-homogeneous results: *RoboLobster* almost always reaches the source when starting at a distance of 50 cm from it, and almost always fails to reach it when starting at a distance of 1 m. Based on these results the authors conclude that the chemical stream is not regular at the periphery of the stream (i.e., at a distance of 1 m from the source), probably due to water turbulence. Now, lobsters significantly outperform the robot in this task. For this reason, the authors conclude that the proposed mechanism is too simplistic—some additional mechanism must be identified to explain how lobsters find the right direction even when the plume is irregular—and that it has to be rejected as it stands. To sum up, in this study, the proposed mechanistic explanation has been discarded because a robotic simulation of it has failed to reproduce the behaviour to be explained.

A second interesting example (referred as “*Plesiosaurus* study” from now on) is reported in (Long 2012; Long et al. 2006). In 1821 a fossil of a giant vertebrate, later called *Plesiosaurus*, was discovered along the coast of the English Channel. The *Plesiosaurus* had four large flippers which were apparently shaped to do the hydrodynamic work of an underwater wing, suggesting that all of them were used to swim. Contemporary living tetrapods using appendages to swim (including penguins and sea turtles) use only either pectoral or pelvic limbs, but not both, for propelling themselves. “Given the species richness of the Mesozoic radiation of tetrapods with apparent four-flippered propulsion, why do more recently derived aquatic tetrapods use only two limbs for propulsion? Which permits better swimming performance, four flippers or two?” (Long et al. 2006, p. 21).

To address this question, John Long and colleagues built a robot, called *Madeleine*, able to swim with four or two flippers. The robot was left free to swim in a pool, using four and two flippers in different experimental sessions, while recording its behaviour in terms of speed, acceleration, and energy consumption. The results indicate that four flippers guarantee more or less the same top speed and acceleration, but—unsurprisingly—they require more energy. From these results specifically concerning *Madeleine*, the authors gained some insights on the adaptive value of four flippers in the *Plesiosaurus*, namely “that the *plesiosaurs*, bearing four putatively-propulsive flippers [...], was likely to have been ambush predators, taking advantage of the better acceleration of four flippers over two” (Long et al. 2006, p. 27). According to the authors, many aquatic tetrapods evolved into two-legged propulsion due to the lower energy consumption occurring in this condition.

To sum up, in this work, an analysis of *Madeleine*'s behaviour informed the authors about the behaviour (speed, energy consumption, acceleration) that the *Plesiosaurus* would have produced in the same conditions. An interesting remark on the relationship between *Madeleine* and the target system, which will be discussed later, is made in Long (2012):

Perhaps the biggest complaint we get about Maddie is that she does not represent any species in particular... That's what we wanted ... I didn't want to

pretend that she was a robotic turtle, for example, as she has come to be named in the popular press. [...] If Madeleine is not a robotic turtle, then how can I claim that she is a robotic *plesiosaurus*? I don't. What I claim is that Maddie uses some of the same propulsive principles that we think both turtles and *plesiosaurus* use and used. Thus, Maddie's mechanistic accuracy is high for any aquatic tetrapod that flaps flippers to swim. Thinking about modelling as the process of representing, Maddie's behaviour represents the behaviour of turtles and *plesiosaurus* in the specific sense that she is about their size and swims with flippers. (p. 188)

## 2.2 Multi-robot Cooperation in Autonomous Mobile Robotics

Autonomous robotics aims to develop robots able to operate in unpredictable environments without continuous human supervision (Siciliano and Khatib 2008). In recent years simulation studies have increased in this field, overcoming skepticism on their actual usefulness. They are usually carried out to analyze the performance of robotic systems or parts of them, and offer several advantages over alternative strategies: they provide a convenient way to explore different robotic scenarios at lower cost with respect to real robots, they reduce the effort needed to write software programs, debug them, and display results, they run faster than their real counterparts. Strong efforts in this community are devoted to proving that the behavior of simulated robots closely matches the behavior of the physical robots themselves: to this end, results obtained in simulations are often—but not always—validated against real experiments.

Some representative examples of simulation studies carried out in autonomous mobile robotics can be found in Balakirsky et al. (2007). Robots can play important roles in disaster mitigation and rescue of human victims. Cooperation among teams of heterogeneous robots, possessing different sensory and motor competencies, is essential in these contexts, mainly due to the variability and unpredictability of the disaster areas (Jennings et al. 1997), where heterogeneous teams combining individual capabilities to solve a task are needed. The Robocup rescue competitions are carried out to boost research on multi-agent rescue robotics and provide benchmarks to evaluate solutions under development. Robocup rescue is structured in two leagues, the rescue robot league and the rescue simulation league, the latter being substantially a competition between simulated rescue robotic teams. One of the divisions of the rescue simulation league is the virtual robots competition, which requires one to simulate small teams of agents with realistic capabilities in small scenarios (other divisions focus on cooperation between larger teams of less realistic simulated robots in larger environments). The establishment of competitions specifically devoted to simulated robotic teams is a signal of the fact that the robotics community is increasingly recognizing the importance of simulations to evaluate the performance of robotics solutions for rescue.

The USARSim simulation software is often used in the virtual robots competitions (Balakirsky et al. 2007). USARSim adds functionalities to the UnrealEngine2 software, which provides a framework to accurately simulate the interaction of

three-dimensional physical objects and enables one to obtain highly realistic simulations of physical systems. USARSim thus allows one to realistically simulate worlds and custom-made robotic systems, carefully modeling sensory inputs and the dynamics of the effectors and the effector-world interaction. In addition, USARSim provides data on the behavior of the simulated robots to evaluate their efficiency. In the virtual robots competitions the simulated world is a disaster scenario and the simulated robots are specifically designed, from a software and hardware perspective, to cooperatively perform the essential tasks involved in rescue, which include the mapping of the environment, the localization of victims, the planning of rescue activities. Two of the following three case-studies make use of this software.

The Virtual Robots 2006 winning team, from the University of Freiburg, Germany, simulated a team of small robots exploiting a novel and powerful idea to perform collaborative environment exploration and self-localization. This simulation study will be referred to as “RFID study” from now on. The basic idea was to let robots release RFID tags (i.e., electronic circuits associated to a worldwide unique number and detectable by suitable devices) at specific points of the explored environment, keeping a trace of those positions in the internal map. From that point on, self-localization in the internal map was aided by the detection of the robots’ relative distance from the RFID tags. The grid of released RFID tags could be used by other robots in the team for self-localization, and as a reference framework to secure consistency among the various local maps of the environment produced by each robot. As described in Kleiner et al. (2006), the efficiency of this solution was tested in several simulation experiments in which simulated robots were left free to explore various virtually reconstructed scenarios and distribute RFID tags, with the simulation environment gathering data about their behavior. These data were then used to assess the simulated robots’ mapping and self-localization ability, which turned out to be fairly good.

Another simulation study—called here the “swarm study”—is described in Cheah et al. (2009). The study focuses on how to coordinate movements of swarms of robots. In the control method proposed there, called region-based shape method, a region is defined within which robots are allowed to move while maintaining a minimum distance among themselves. The desired region can be specified as an arbitrary shape. The robots in the group only need to communicate with their neighbors and do not have any specific identity or role in the group. Therefore, the proposed method does not require one to impose any specific order or position to the robots, and different shapes can be formed by the group. The performance of the region-based shape controller is tested in a simulation comprising of 100 robots moving along a path. The actual mass of each robot is set at 1 kg and the actual value of the damping constants (describing how oscillations in a system decay after a disturbance) is set to 0.5 kg. The desired minimum distance among robots is set to 0.3 m. The robots proved able to move into regions shaped as a circle, as a ring, and as a geometric figure defined by arbitrary functions.

In other studies, the output of the simulation is compared with the behavior of real-life robots operating in the same circumstances and engaging in the same task, a comparison which was not carried out in the two previously described studies. Another participating team at the 2006 Virtual Agents competition, from the

University of Pittsburgh, developed an algorithm—called *Steel*—to coordinate teams of rescue robots with human intervention (this study will be called “Steel study” from now on). In the Steel study the robotic system is not fully autonomous, as the human user could intervene at various levels, e.g., by changing robots’ plans, forcing them to regenerate a plan, reconfiguring goal priorities. Experiments have been carried out to test the efficiency of the system. Participants controlled real-life and simulated robots in different experimental sessions, using the same interface. In both cases the task was to guide the robot along a narrow corridor with varying types of debris (wood floor, scattered papers, lava rocks), following a straight or a complex path, and avoiding obstacles, either teleoperating the robot or controlling more high-level parameters of its behavior. Task completion time, distribution of commands and pauses, changes in heading, and number of issued commands were reported to be similar in the two conditions.

A summary of the biorobotics and autonomous mobile robotics case-studies discussed so far is provided in Table 1.

### 3 Key Concepts and Methodological Commonalities

#### 3.1 Theoretical Models

The studies described in Sects. 2.1 and 2.2 belong to very different research traditions and communities in robotics. Yet, despite their surface differences, they share a common methodological structure: in all of them, an artificial system is used to acquire knowledge about a target system, either natural or artificial. The artificial system simulates the target system under a certain theoretical perspective. Some clarification on how key concepts such as “simulation” and “theoretical model” are used here will be useful to outline this common methodological structure.

In simulation studies, a simulation system is used to acquire knowledge about the behaviour of a target system. This is the case of the lobster study, where the simulation system is *RoboLobster* and the target system is any individual of the species of lobsters. In the *Plesiosaurus* study, the simulation system is *Madeleine* and the target system is any aquatic tetrapod that flaps flippers to swim. In two of the three autonomous robotics studies described in Sect. 2.2 (namely the RFID and Steel studies), the simulation systems are programmed computers<sup>2</sup> running the USARSim framework as simulation environment and the target systems are (teams of) real-life robots. As more extensively discussed later, some uses of the simulation system require the target system to be a real-life system, while other uses do not. Indeed, in biorobotics, the simulation system is a robotic system and the target system is

---

<sup>2</sup> The question whether the hardware or the software of a computer is what simulates the target system (or a model of it) in computer simulation studies has been debated in the philosophical literature (Barberousse et al. 2009; Beisbart 2018). Here the simulation system is taken to be a programmed computer described symbolically, i.e., in terms of variables taking values and of relationships holding among them. However, different interpretations of the notion of “programmed computer” may be compatible with the analysis made here.

**Table 1** A list of the case-studies analysed in this article

	Study	Description	Bibliographic reference
Biorobotics	Lobster study	An underwater robot called RoboLobster is used to test a mechanistic hypothesis on the chemotaxis of lobsters	Grasso et al. (2000)
Autonomous mobile robotics	<i>Plesiosaurus</i> study	An underwater robot called Madeleine is used to study the swimming behaviour of an extinct animal, the <i>Plesiosaurus</i>	Long (2012) and Long et al. (2006)
	RFID study	A computer simulation of rescue robots releasing RFID tags in the environment for self-localization and mapping provide insights on the efficiency of the algorithm	Kleiner et al. (2006) and Balakirsky et al. (2007)
	Swarm study	A region-based shape method for coordinating movements of swarms of robots is tested in simulation	Cheah et al. (2009)
	Steel study	A computer simulation of a rescue robot is implemented to test Steel, an algorithm for coordinating teams of rescue robots with human intervention. The behaviour of the simulation is compared with the behaviour of a real-life robot running Steel	Balakirsky et al. (2007)






Fig. 1 The relationship between simulation systems, theoretical models, target systems

an existing living system. In the steel study, the target system is an existing robot and the simulation system is a programmed computer. However, in two of the three autonomous robotics studies described before—the RFID and the swarm study—the target system is a fictional entity. No real-life counterpart of the simulated RFID-releasing multi-robot team is reported to exist in Balakirsky et al. (2007), and no real-life counterpart of the simulated 100 robots team is reported to exist in Cheah et al. (2009).

The expression “simulation system” is used to refer to the system which is assigned the role of simulating something. What is the “something” being simulated? Statements of the form “artificial system A simulates target system S” are quite common in scientific parlance. However, as often remarked in the philosophical literature, simulation systems simulate the target system under a certain theoretical perspective. There is no such thing as a simulation of a lobster. Many simulation systems can be said to simulate a lobster, differing from one another in the way they represent lobsters. *RoboLobster*, for example, simulates lobsters under a certain theoretical perspective according to which each chemoreceptor stimulates the contralateral motor organs. Similarly, *Madeleine* simulates aquatic tetrapods under a particular theoretical perspective which abstracts away from a number of characteristics of them, including the species, while representing size and number of organs for propulsion (recall John Long’s quotation reported in Sect. 2.1). For this reason, in what follows it will be assumed that simulation systems simulate theoretical models of other system and not “directly” those systems (see Fig. 1), even though the shorter formulation will be occasionally used for the sake of brevity.

“Theoretical models” (or models, from now on) are taken here to be non-concrete interpreted structures (Weisberg 2013, chapter 3) characterizing the target system in terms of parameters which may take values (see also Winsberg 1999). Parameters taking certain values are interpreted as property ascriptions (see Weisberg 2013, chapter 3 for a discussion on the interpretation of non-concrete models). Theoretical models also prescribe that certain regularities hold among these parameters: these regularities are interpreted as regularities holding among properties of the target system.<sup>3</sup> Theoretical models are not to be confused with their descriptions, a model being typically describable in many ways, for instance in mathematical or computational terms. Theoretical models are dynamic in the sense expressed by Hartmann

<sup>3</sup> This construal of the concept of “theoretical model” is substantially akin to the analysis of models as set-theoretic structures (Suppe 1989).

(1996) if they establish relationships between the current value of some parameters and future values of the same or of other parameters.

The theoretical model simulated in the lobster study, for example, characterizes the target system in terms of parameters representing the intensity of left and right chemoreceptive stimuli and the activity of left and right motor organs, stating relationships between the two. The theoretical model simulated in the *Plesiosaurus* study characterizes aquatic tetrapods in terms of their size and some motor parameters defining the number of flippers and the gait. In both studies, the structure of the theoretical model is explicitly stated and discussed by the authors of the study. In other cases, it is not. This does not imply, however, that the simulation system does not simulate a theoretical model of the target system: it only implies that the theoretical model has not been explicitly stated. Indeed, in the RFID, swarm, and Steel studies, the computer implements a theoretical model of the (fictional or existing) target robots. In the RFID study, the computer simulates a very detailed theoretical model of the real-life robot *Zerg* which “captures the same physical properties as the real one, e.g. a four wheel drive, a RFID tag release device, a RFID antenna, inertial measurement unit (IMU), and LRF [...]. The sensors of the model are simulated with the same parameters as the real sensors, expect the real RFID reading and writing range” (Balakirsky et al. 2007, p. 11). The control algorithm is part of the theoretical model, as it establishes relationships between sensor readings, motor outputs, and the value of internal parameters which, in the computer simulation of this model, will be implemented as portions of system memory. In the swarm study a much less detailed, point-like theoretical model of a robot is simulated instead. The theoretical model includes the dynamics equation modelling the behaviour of each robot.

The expressions “how-actually” and “how-possibly” are often used to characterize the relationship between theoretical models and target systems: if *M* is a how-actually model of target system *S*, then *M* correctly represents the relationships holding between some of *S*' properties, while *M* is “how-possibly” if it is a loosely constrained conjecture about *S* (Craver 2006). Following Glennan (2017), it will be assumed here that “there is nothing that intrinsically distinguishes a how-possibly from a how-actually model. The difference rests solely on whether or not the theoretical hypotheses concerning the model-target relation are known to be correct” (p. 69).

Theoretical models can be used to predict or explain the behaviour of the target system. In the lobster study the theoretical model of *RoboLobster* is intended to provide a basis to explain the chemotactic behaviour of lobsters, by identifying the underlying mechanism.<sup>4</sup> On the contrary, the RFID study has no explanatory purpose: rather, a theoretical model of RFID-like real-life robots is simulated in a machine to predict the behaviour that real-life *Zerg* robots would produce in the same circumstances—more specifically, to predict whether they would be able to

<sup>4</sup> It will not be assumed that all theoretical models are explanatory. The problem of what makes a theoretical model explanatory is out of the scope of this article: for an up-to-date discussion, see Bokulich (2017).

complete a given rescue task. The point at stake here is that, in both cases, what is under investigation is the behaviour of the target system. In what follows it will be assumed that target system behaviour can consist in the values taken by some system parameter at a certain time, or in the trajectory of these values over a certain period. For example, lobster's behaviour can be represented as the trajectory of the values of motor organs activities over a specific period. Different parameters can be chosen to define target system behaviour, depending on the theoretical interests of the experimenter. The USARSim environment, as pointed out before, enables one to monitor and record the value taken by those parameters in time, thus allowing the experimenter to obtain information about the behaviour of the simulation system.

In all the studies considered here, the behaviour of the simulation system is observed in a particular environment and during the execution of a particular task (e.g., a swimming task in the biorobotics examples, a rescue task in the autonomous mobile robotics tasks). The environment is a real-life one in the biorobotics case-studies and a simulated one in the autonomous mobile robotics studies. In biorobotics there is an evident physical distinction between the robot and the environment, whereas, in computer simulations of robotic systems, the very same programmed computer may run a simulation of (a theoretical model of) the target system and of the environment. However, the two are to be kept conceptually distinct from one another. Indeed, one can typically intervene on the structure of the simulation system—i.e., on the specific part of the programmed computer which implements the theoretical model of the target system—without intervening on the structure of the environment—i.e., on the part of the programmed computer implementing a simulation of the environment. This is often done to assess whether changes in the robot's structure or control system guarantee better performances in the same environment, and whether the same robot can perform equally well in different environments. Note also that, in autonomous mobile robotics, the simulated environment is typically a less detailed and simplified version of the corresponding real-life environment. This does not fundamentally differ, however, from what happens in biorobotics. *RoboLobster* and *Madeleine* (in the lobster and *Plesiosaurus* studies) are put in carefully designed laboratory environments which are significantly different from those in which real-life lobsters and aquatic tetrapods live. Thus, as discussed in Tamburrini and Datteri (2005) in connection with biorobotics, in both cases extra assumptions and arguments are needed to generalize the results obtained in simulation to what would happen in a real-life environment.

### 3.2 The Relationship Between Theoretical Models and Simulation Systems

The term “simulation” enjoys a broad range of uses in the scientific and philosophical literature. In some cases, it refers to an act or process (e.g., “simulations were performed of layers I and II of olfactory paleocortex”, Ambros-Ingerson et al. 1990). In other cases, it refers to a system (e.g., “In total, our simulation comprised over 1800 cells and 6000 synapses”, Blanchard et al. 2000). Something can be a simulation of a (biological) system, of a behaviour (Feigenbaum 1961), or of a process (Simon and Newell 1962). Here, as pointed out before, the expression “simulation

system” is used to refer to the system which has the role of simulating a theoretical model of the target system. Depending on the context, the term “simulation” is used here to mean the act of running a computer or a robotic simulation system—i.e., of turning it on and observing its behaviour—or, for the sake of brevity, to denote a simulation system.

A central problem in the philosophy of simulation literature, re-phrased in the terminology adopted here, is to identify the conditions under which a simulation system can be said to simulate a theoretical model of the target system. This problem is relevant for the present discussion, a central claim of this article being that some apparently distant areas of robotics research are often concerned with simulating theoretical models of existing or non-existing systems. Several answers to this question have been proposed in the literature (Hartmann 1996; Humphreys 2004). Here the following working definition will be adopted. As pointed out before, theoretical models characterize target systems in terms of parameters and regularities holding among them. System A simulates theoretical model M only if A can be characterized in terms of parameters whose values depend on one another according to the regularities mentioned in M. For example, suppose M is the theoretical model formulated in the framework of the lobster study. To build a computer simulation of that model one writes a program, in a suitable programming language, defining variables representing left and right chemotactic intensity and the activity of the motor organs. Pieces of code will establish the appropriate relationships between the values of these variables.

Note that, in computer and cognitive science, the term “simulation” is typically used to refer to cases in which the simulation system is a suitably programmed computer. Interestingly, some authors—e.g., Ziemke (2003)—use that term in a way that implies that robotic models of living system behaviours should not be properly regarded as simulations of anything, exactly because they are robotic, embodied systems and not programmed computers. From this view it would follow that a system such as *RoboLobster* in the lobster study—which is a real-life robot and not (only) a programmed computer—should not be properly called a simulation (of a theoretical model of lobster behaviour). One should be careful to note, however, that the relationship between *RoboLobster* and the theoretical model of lobster chemotaxis is similar, at a certain level of analysis, to the relationship which would hold between the latter and the programmed computer mentioned in the previous paragraph. Indeed, *RoboLobster* can be characterized in terms of parameters whose values depend on one another according to the regularities mentioned in that model. A stable relationship holds between the pattern of electrical signals detectable at the output of the left chemotactic sensor mounted on the robot and the pattern of electrical signals delivered to the contralateral motors, a relationship which matches the corresponding relationship defined in the model. For this reason, *RoboLobster* and *Madeleine* are regarded here as robotic systems simulating theoretical models of lobsters and aquatic tetrapods.

Three remarks on the notion of a simulation are needed. First, the question whether simulation system A simulates the theoretical model M is orthogonal to the question whether M describes how the target system S works. One can accurately simulate a “wrong” model of S, and one can inaccurately simulate a how-actually

model of S. Second, the definition provided here only establishes necessary conditions for qualifying A as a simulation of M: additional conditions may be needed, whose identification is out of the scope of this paper.

Third, in simulation studies the simulation system is used to obtain the behavioural implications of the simulated theoretical model. The behaviour of the simulation system is interpreted as the behaviour that the target system would generate if the theoretical model implemented in the machine was a how-actually model of it. For example, the behaviour of *RoboLobster* in the lobster study is interpreted as the behaviour that a lobster would generate, in the same conditions, if the chemotactic theoretical model implemented in the machine was a how-actually model of lobster chemotaxis. The behaviour of the RFID simulation system in the RFID study is interpreted as the behaviour that a real-life *Zerg* robot, having the hardware and control characteristics specified in the theoretical model, would generate in the same conditions. Note the hypothetical structure of this claim: it does not imply that, in simulation analyses, one assumes that the simulated theoretical model is how-actually. The main role of a simulation system is to generate the behavioural implications of the implemented theoretical model, independently of whether the latter is a plausible model of the target system or not.

#### 4 Some Peculiarities of Biorobotics and Autonomous Robotics Simulations

The case-studies described in Sects. 2.1 and 2.2 belong to relatively distant provinces of robotics. However, they share a common methodological structure, which can be summarized here using the terms discussed in the previous section. They all involve a simulation system, which simulates a theoretical model of another system. The simulation system is used to obtain the behavioural implications of the simulated theoretical model, i.e., to produce the behaviour that the target system would produce in the same circumstances if the model was how-actually. In all cases, the behaviour of the simulation system is observed in an experimental environment (which is a suitably constrained real-life laboratory setting in the biorobotics cases, and a computer simulated environment in the autonomous mobile robotics cases) during the execution of a certain task, which may involve finding a chemical source, swimming, or rescuing people in a disaster scenario. Given this common structure, the two classes of studies described before differ from one another in a number of interesting respects.

A first surface-level difference between the biorobotics and autonomous robotics case-studies concerns the role of the robotic system in this common structure. In the biorobotics studies, the robotic system plays the role of simulation system. In the autonomous mobile robotics ones, the robotic system is the target system itself (and the simulation system is a programmed computer). A second difference concerns the role of the simulation system with respect to the behaviour of the target system. In the biorobotics case-studies analysed before, the simulation system is used for explanatory purposes, i.e., to test the plausibility of the simulated theoretical model. In the autonomous robotics case-studies, the simulation system is used for

predictive purposes, i.e., to predict the behaviour of rescue robots. This distinction can be brought to bear on the fact that, in the second class of simulation studies only, the target system can be fictional.

Indeed, as occasionally pointed out in the philosophical literature (e.g., Guala 2002), simulation systems can be used at least for two kinds of epistemic purposes. In some studies, which are called “model-oriented” here, the behaviour of the simulation system A is compared to the behaviour of the target system S. If the two behaviours match, one may be induced to infer—under some auxiliary assumptions—that M can be included in the space of the how-possibly theoretical models of S’ behaviour.<sup>5</sup> On the contrary, if the two behaviours do not match, one may be induced to exclude the model from that space. Model-oriented studies are carried out to answer questions of the form “How does S produce the behaviour under investigation?”. In case of behavioural match, the “new knowledge” ultimately gained about S consists in the description of a how-possibly theoretical model of the target behaviour.

In other studies, called “prediction-oriented” here, simulation system A’s behaviour is interpreted as the behaviour that the target system S would actually produce under the same conditions. This interpretive step is not supported by comparisons between A’s and S’ behaviours, as in the model-oriented strategy. Indeed, prediction-oriented studies are often carried out when independent information on A’s output is difficult or impossible to obtain for practical or theoretical reasons. Prediction-oriented simulation studies are also carried out when direct access to S’ behaviour would be more expensive or time consuming. They address questions of the form “How will S behave?” and the “new knowledge” obtained about S, contrary to the model-oriented case, consists in a description of S’ behaviour. In prediction-oriented (but not in model-oriented) studies, the simulation system A is used as a sort of surrogate of S, i.e., as a system on which one is supposedly legitimated to make experiments and perform measurements *as if it was S* to gain information on its real-life behaviour (see also Frigg and Nguyen 2017; Swoyer 1991 on model-based surrogative reasoning). The interpretive step characterizing prediction-oriented simulations is not taken in the model-oriented strategy. In a model-oriented strategy, one has no reason to believe that A’s behaviour will match S’ behaviour: the comparison will tell.

The model-oriented strategy is aptly exemplified by the lobster study, whose purpose is to test a theoretical model of lobster chemotaxis. The *Madeleine* robot in the *Plesiosaurus* study, on the contrary, was not used to test a model of tetrapod swimming, but to obtain the behaviour that a tetrapod of roughly the same size of a *Plesiosaurus* would generate (in terms of velocity, acceleration, and energy consumption) using four flippers instead of two. The vast majority of biorobotics studies are model-oriented: indeed, the *Plesiosaurus* study is only a notable exception. The

<sup>5</sup> The model-oriented strategy has been sometimes called “synthetic method” in artificial intelligence and cybernetics (Cordeschi 2002). Note that this strategy can only lead one to reduce or increase the space of the how-possibly theoretical models of S’ behaviour. A’s reproduction of the latter, per se, guarantees neither that M is the only possible model of it nor that it is explanatory.

reason may be easily understood. As pointed out before, prediction-oriented studies are carried out to obtain behavioural information that cannot be easily accessible otherwise for theoretical or practical reasons, i.e., because no adequate measurement tool is available—for example, one simulates the conformational changes of ion channels (Dror et al. 2012) because no technique is currently available to obtain them at the same resolution and level of detail—or because the target behaviour occurred in the past or will occur in the future—as in meteorological simulations for the production of weather forecasts. The *Plesiosaurus* study, as a matter of fact, concerned the behaviour of an extinct animal, being in principle inaccessible. However, biorobotic simulation studies are typically carried out to investigate on the behaviour of living organisms, i.e., behaviours which are in most cases accessible through conventional experimental tools. There is little need of prediction-oriented studies in these cases.

To be sure, what characterizes the model-oriented strategy is not the mere fact that A's and S' behaviour are compared, but the fact that the outcome of the comparison is used as a premise to shape the space of the how-possibly models of S' behaviour. On the one hand, the mere absence of a comparison is not a good reason to conclude that the analysis in question is prediction-oriented: it may be simply the case that data on S' behaviour are yet to be acquired. On the other hand, machine outputs can be compared with S' behaviours for purposes other than theoretical model corroboration or rejection.

In autonomous mobile robotics, on the other side, computer-based simulation systems are used both for predictive purposes, i.e., to obtain information on the target robotic system before building it (information that is currently inaccessible, as the system is yet to be built) and for model-oriented purposes, i.e., to understand the relationship between the theoretical model and the target system. In the first case the target system can be fictional, as there is no need to compare the behaviour of the simulation and the target system, given that the purpose is predictive. This is evident in both the RFID study and the swarm study where simulated robots are not real-life robots.<sup>6</sup> In the second case the target system is required to be existing and accessible, as a key element of the model-oriented strategies is the comparison between the behaviour of the simulation and the target system. The authors of the Steel study, for example, compared the performances of simulated and real-life rescue robots to understand whether the simulated robotic system is able to correctly simulate the target system, i.e., to check whether the simulation system incorporates a good theoretical model of the target system.

---

<sup>6</sup> An interesting difference between these two simulation case-studies is worth emphasizing here. In the RFID study the simulated robot is based on the real-life robot *Zerg* and reproduces the same physical properties of the real-life one. The robots simulated in the swarm study, on the contrary, are purely fictional and not modelled after any real-life robots. Still, both studies have predictive purposes. The purpose of the swarm study, in particular, is to predict the behavior of purely fictional entities which reproduce no existing robot.



## 5 Concluding Remarks

In this paper two broad classes of simulation studies in robotics have been identified and discussed in order to shed light on the epistemic role of simulations in robotics. In the first class, simulation studies aim to acquire knowledge on living systems. In the second class, simulation studies are carried out to acquire knowledge on the performances of robotics systems. These two classes of studies have been shown to have a common methodological structure: in particular, it has been argued that both classes involve simulations of theoretical models of target system. The notion of “robotic simulation” has been introduced and discussed to reflect on the first class of studies, illustrated by a couple of biorobotics case-studies.

The distinction between model-oriented studies and prediction-oriented studies has been introduced to discuss some methodological peculiarities of each class of simulation studies. Whether the differences emerged in Sect. 4 are to be generalized beyond the case-studies of this paper remains an open question which is worth analyzing in future research. However, the common methodological structure proposed here for robotic simulations and computer simulations of robots can pave the way to revealing further methodological peculiarities in biorobotics and autonomous robotics simulation studies.

## References

- Ambros-Ingerson, J., Granger, R., & Lynch, G. (1990). Simulation of paleocortex performs hierarchical clustering. *Science*, *247*(4948), 1344–1348.
- Arkin, R. C. (1998). *Behavior-based robotics*. Cambridge: The MIT Press.
- Balakisrsky, S., Carpin, S., Kleiner, A., Lewis, M., Visser, A., Wang, J., et al. (2007). Towards heterogeneous robot teams for disaster mitigation: Results and performance metrics from Robocup rescue. *Journal of Field Robotics*, *24*(11), 943–967.
- Barberousse, A., Franceschelli, S., & Imbert, C. (2009). Computer simulations as experiments. *Synthese*, *169*(3), 557–574.
- Beisbart, C. (2018). Are computer simulations experiments? And if not, how are they related to each other? *European Journal for Philosophy of Science*, *8*(2), 171–204.
- Blanchard, M., Rind, F. C., & Verschure, P. F. M. J. (2000). Collision avoidance using a model of the locust LGMD neuron. *Robotics and Autonomous Systems*, *30*(1–2), 17–38.
- Bokulich, A. (2017). Models and Explanation. In L. Magnani & T. Bertolotti (Eds.), *Springer handbook of model-based science* (pp. 103–118). New York and Cham: Springer.
- Braitenberg, V. (1986). *Vehicles. Experiments in synthetic psychology*. Cambridge: The MIT Press.
- Cheah, C. C., Hou, S. P., & Slotine, J. J. (2009). Region-based shape control for a swarm of robots. *Automatica*, *45*(10), 2406–2411.
- Cordeschi, R. (2002). *The discovery of the artificial. Behavior, mind and machines before and beyond cybernetics*. Dordrecht: Springer.
- Craver, C. F. (2006). When mechanistic models explain. *Synthese*, *153*(3), 355–376.
- Datteri, E. (2017). The epistemic value of brain–machine systems for the study of the brain. *Minds and Machines*, *27*(2), 287–313.
- Datteri, E., & Tamburrini, G. (2007). Biorobotic experiments for the discovery of biological mechanisms. *Philosophy of Science*, *74*(3), 409–430.
- Dror, R. O., Dirks, R. M., Grossman, J. P., Xu, H., & Shaw, D. E. (2012). Biomolecular simulation: A computational microscope for molecular biology. *Annual Review of Biophysics*, *41*(1), 429–452.



- Feigenbaum, E. (1961). The simulation of verbal learning behavior. In *Papers presented at the May 9–11, 1961, Western joint IRE-AIEE-ACM computer conference*, pp. 121–132.
- Frigg, R., & Nguyen, J. (2017). Models and representation. *Springer handbook of model-based science* (pp. 49–102). Cham: Springer.
- Glennan, S. (2017). *The new mechanical philosophy*. Oxford: Oxford University Press.
- Grasso, F. W., Consi, T. R., Mountain, D. C., & Atema, J. (2000). Biomimetic robot lobster performs chemo-orientation in turbulence using a pair of spatially separated sensors: Progress and challenges. *Robotics and Autonomous Systems*, 30(1–2), 115–131.
- Guala, F. (2002). Models, Simulations, and Experiments. In L. Magnani & N. J. Nersessian (Eds.), *Model-based reasoning. Science, technology, values* (pp. 59–74). New York: Springer US.
- Hartmann, S. (1996). The world as a process: simulations in the natural and social sciences. In R. Hegselmann, et al. (Eds.), *Simulation and modeling in the social sciences from the philosophy of science point of view* (pp. 77–100). Dordrecht: Theory and Decision Library, Kluwer.
- Humphreys, P. (2004). *Extending ourselves: Computational science, empiricism, and scientific method*. New York: Oxford University Press.
- Jennings, J. S., Orleans, N., Whelan, G., & Evans, W. F. (1997). Cooperative search and rescue with a team of mobile robots. In *Proceedings of the IEEE Int. Conf. on Advanced Robotics (ICAR)*, (pp. 193–200).
- Kleiner, A., Prediger, J., & Nebel, B. (2006). RFID Technology-based exploration and SLAM for search and rescue. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, (pp. 4054–4059).
- Long, J. (2012). *Darwin's Devices. What evolving robots can teach us about the history of life and the future of technology*. New York: Basic Books.
- Long, J. H., Schumacher, J., Livingston, N., & Kemp, M. (2006). Four flippers or two? Tetrapodal swimming with an aquatic robot. *Bioinspiration & Biomimetics*, 1(1), 20–29.
- Pfeifer, R. (2009). Biologically inspired robotics. *Science*, 1088 (2007).
- Rosenblueth, A., & Wiener, N. (1945). The role of models in science. *Philosophy of Science*, 12(4), 316–321.
- Siciliano, B., & Khatib, O. (Eds.). (2008). *Springer handbook of robotics*. Heidelberg: Springer.
- Simon, H. A., & Newell, A. (1962). Computer simulation of human thinking and problem solving. *Mono-graphs of the Society for Research in Child Development*, 27(2), 137.
- Suppe, F. (1989). *The semantic conception of theories and scientific realism*. Urbana and Chicago: University of Illinois Press.
- Swoyer, C. (1991). Structural representation and surrogative reasoning. *Synthese*, 87(3), 449–508.
- Tamburrini, G., & Datteri, E. (2005). Machine experiments and theoretical modelling: From cybernetic methodology to neuro-robotics. *Minds and Machines*, 15(3–4), 335–358.
- Webb, B. (2001). Can robots make good models of biological behaviour? *The Behavioral and Brain Sciences*, 24(6), 1033–1050–1094.
- Webb, B. (2006). Validating biorobotic models. *Journal of Neural Engineering*, 3, R25–R35.
- Weisberg, M. (2013). *Simulation and similarity. using models to understand the world*. Oxford: Oxford University Press.
- Winsberg, E. (1999). Sanctioning models: The epistemology of simulation. *Science in Context*, 12(2), 275–292.
- Ziemke, T. (2003). On the role of robot simulations in embodied cognitive science. *AISB Journal*, 1(4), 389–399.