Part II

Problem-Solving Robots

The preceding part presented human problem solvers who have been classified into several intersecting categories, each one exemplifying a specific method for attacking problems. The chapters were conceived as a woven fabric interlacing historical problem-solving in science and technology with robotics projects elaborated by young enthusiasts attending the LEGO MINDSTORMS robotics lab called *B.I.T.*, which is a modest research center for youths at the periphery of our boarding institution. Through the patchwork-like pages, readers were introduced to divergent ideas, facts, and knowledge from different research fields, elucidating robotics-related aspects of mechanics, electronics, computer sciences, and mathematics. They were challenged to solve relevant problems in order to train their skills in the various disciplines and to discover their own very personal approaches to solving problems. Many exercises concerned robotic automata, others machines, or computers, and some explored more basic subjects from electronics and mechanics. It was noted that the distinctions between robots, automata, and machines are not as evident as one might think, and that definitions of these devices are often incoherent and inconsistent, depending entirely on the definer’s point of view.

Chapters 9 through 11 will explore some aspects of the vast field of robotics from a problem-solving perspective. Because the coming examples, exercises, and challenges are exclusively based on real LEGO robots, we purposely stick to the definition of a robot offered by Maja J. Matarić[1], stating that “a robot is an autonomous system which exists in the physical world, can sense its environment, and act on it to achieve some goals” [1]. This description includes many items in my own house: the bread baking machine, washer, dryer, freezer, DVD-recorder, and our combined heating and hot water supply system—although nobody would call them robots. The tiny iRobot vacuum cleaner and autonomous lawnmower are immediately perceived as robots, probably because they are mobile, which provides them with some familiar attributes of animal life or animation. As all of the robots presented in this section will show

[1] Cited in Chapter 5.

these attributes of animation, we will refine the definition by stating that "a robot is an autonomous system, which **moves** in the physical world, can sense its environment, and act on it to achieve some goals." Note that the restriction is only apparent because an immobile robot can be viewed as a robot with zero motion if we include the absence of anything as a concept extension, as did the lonely Hindu inventor of the number zero.

The ability to solve problems is not limited to humans but is a fundamental capacity of every living being that is submitted to existential issues of nurturing, reproduction, nesting, and defense—all of which require some form of intelligence that must be seen as the *basic prerequisite* of problem solving. Unfortunately there is no consensus on the definition of the term **intelligence**. On the contrary, its meaning substantially diverges from the function of the invoked disciplines. Human subject research describes intelligence as a compound faculty of higher adaptation and reasoning, including learning from experience or insight, abstract thinking, communication (by speech, mimic, and gesture), anticipation, planning, use of symbolism, creativity, and problem solving, which are necessarily based on a highly efficient brain. From this point of view, intelligence is considered in relation to the fine motor skills of the human hand, the invention of tools, the self-concept and consciousness of the individual, and his or her social interactions. In order to mark the boundary between the human and animal worlds, the term **animal cognition** is preferred to the word **intelligence** if related animal skills are being depicted and eventually compared to their human counterparts.

By 1637, the great French philosopher and mathematician René Descartes had compared the animal with a banal automaton as being incapable of thought. However ingenious and complex it may appear, he reasoned, it is nothing more than "a machine made by the hands of God, which is incomparably better arranged, and adequate to movements more admirable than is any machine of human invention" [2]. Captured in a pronounced dualism and convinced of the immateriality of the human mind, he argued, "that not only the brutes have less reason than men, but that they have none at all." It was inconceivable for him that an animal or a machine could think. Anticipating the discussion about intelligent machines, he proposed two imperturbable means of distinguishing the irrational creature from the human. First, he considered speech to be the expression *par excellence* of thought. Even if an automaton could be made capable of talking, it would just emit words, perhaps rattle on, but never give an appropriate answer to anything being said in its presence as the most unintelligent person can do. Second, Descartes stated that the actions deployed by the animal or machine, however skillful they might appear, would never be the result of knowledge or insight but would be just the dummy product of ingenious "organs."

In the discussion about *Maelzel's Chess Player* (see Challenge 7, Chapter 3), we saw that the idea of an automaton capable of thinking and acting in a human manner intrigued scientists, artists, and poets for almost a century. However, the preoccupation with this subject was already reported from antiquity, since automata were used in ancient Greece, Anatolia, and China. Hephaestus, who is the Greek god of technology, blacksmiths, artisans, and sculptors, is said to have built bronze and golden human machines [3]. In *The Iliad*, Homer narrated:

> "Huge god Hephaestus got up from the anvil block with laboured breathing. He was lame, but his thin legs moved quickly under him... At once he was helped along by female servants made of gold, who moved to him. They looked like living servant girls, possessing minds, hearts with intelligence, vocal chords, and strength. They learned to work from the immortal gods. These women served to give their master detailed help." [4]

This same concept repeatedly appeared in variations as the *Pygmalion* myth and has been treated in numerous stories, starting with diverse interpretations of *Le Avventure Di Pinocchio: Storia Di Un Burattino* by Carlo Colodi (1881), Fritz Lang's *Metropolis* (1927), to Steven Spielberg's *A.I. Artificial Intelligence* (2001). When the word **intelligence** was first used in the context of machines, there was (and still is) a great controversy surrounding the concept. In 1956, a summer workshop held at Dartmouth University in Hannover, New Hampshire, gave birth to the new field of **artificial intelligence (AI)** [5]. One of the prominent attendees of that conference, John McCarthy (1927–2011) late professor of Computer Science at Stanford University, and he defined artificial intelligence (AI) as "the science and engineering of making intelligent machines, especially intelligent computer programs." He related AI to the task of using computers to understand human intelligence, which he considered to be the computational part of the faculty to achieve goals in the world [6]. Much of the early optimistic work of AI consisted of trying to emulate human thinking by the means of computers. Alan Turing, who in 1950 had already discussed the possibilities and the difficulties of machine intelligence in his pioneering paper "Computing Machinery and Intelligence" [7], provided an astute method of finding out whether a machine can think. The famous **Turing test**, which is a descendent of Descartes' method for distinguishing between the irrational creature and the human, was presented as follows:

1. An interrogator puts questions to two participants, who are seated in separate rooms.
2. The participants type their responses.
3. The interrogator can neither see nor hear the participants, and is able only to read their typed responses.
4. One participant is human, while the other is a machine engaged in a human "imitation game," programmed to respond to the questions as a human would.
5. The interrogator must determine which of the two participants is the machine.
6. Ultimately, if the interrogator cannot discern the difference, then the machine should be judged capable of human thought.

So far, no machine has passed this test, despite promising achievements with masterpieces of programming like expert systems, computer-developed mathematical proofs, neural networks, and the incarnation of the Chess Player called *Deep Blue*©, which lost the first but won the second of two runs against World Champion Garry Kasparov in 1996 and 1997. In an auto-critical evaluation [8], IBM explains that *Deep*

Blue, which is a powerful supercomputer, could analyze up to 200,000,000 chess positions per second, compared to Kasparov's "poor" three chess positions per second. By contrast, the grandmaster relied on a vast knowledge of moves and gambits, surmounting *Deep Blue's* limited database. Garry Kasparov, who is a world-class player, deployed his extraordinary intuition combined with learning and adaptation capacities, of which *Deep Blue* was totally incapable. *Deep Blue* was not a "learning system"; it didn't use artificial intelligence to learn from its adversary or "think" about the current constellation of the pawns. Finally, Deep Blue was entirely dependent on five IBM research scientists and one international grandmaster, who adapted the program during the matches.

The experience with *Deep Blue* typifies the general sense of disillusion experienced during the attempts to mimic human intelligence. In 1991, Rodney A. Brooks, who is emeritus Panasonic Professor of Robotics at MIT, noted in his ruthlessly candid paper "Intelligence without representation" [9] that AI started as a domain whose goal was the replication of human intelligence in a machine. However, early hopes vanished as the dimension and difficulty of the task became obvious, so instead the research retreated into specialized subproblems, such as knowledge representation, natural language understanding, vision, and so forth [7]. Despite the enormous progress made through the spectacular increase in computational power, the greatest difficulties came up when AI researchers ventured into complex problems that provoked an exponential growth in the memory size and the computing time required. It seemed that many AI experts had fallen into Sessa's trap and had to learn a sour lesson about modesty (see Chapter 6, Glossary: Exponential Growth). More specifically, AI roboticists understood that even simple human cognitive capacities and related activities are much more difficult to implement into a machine than anyone thought. Obstacles arose at almost every corner of the research in relation to the fundamental understanding of problem solving based on effective heuristics, planning, insight, creativity, and internal models of the world—going beyond simple pigeonholing of collected information and applying astute search algorithms.

Then came Watson®, IBM's novel supercomputer, which is described as the world's most advanced question-answering system. In February of 2011, this computer accessed its extraordinary knowledge base and challenged human *Jeopardy!*[2] all-star champions Ken Jennings and Brad Rutter in a three-night game... and won! Before the victory, probably hesitant about the outcome of the challenge, IBM humbly declared that Watson had developed from a modest DeepQA machine to a formidable *Jeopardy!* contestant [11]. The superbrain, made up of 90 servers, applies advanced natural language processing, ultra-fast information retrieval, AI-based reasoning, and machine learning technologies to open-domain question answering. IBM compares its "youngest child" with the science-fictional Star Trek onboard computer, which is an interactive dialog agent that can answer questions and provide precise information on any topic [12]. Capable of expressing itself in human terms, Watson is perhaps the first machine to come within reach of passing the Turing test.

[2] *Jeopardy!* is a popular American quiz show, running from 1964 to this day, that requires contestants to formulate their responses in question form.

However, Dr. Chandra, the fictive creator of Arthur C. Clarke and Stanley Kubrik's *HAL9000*,[3] along with "the dreamers still in the field of AI" [13], will have to wait before eventually realizing the first machine capable of proclaiming Descartes's *Cogito ergo sum*[4] and philosophizing "But what then am I? A thing that thinks. What is that? A thing that doubts, understands, affirms, denies, wills, refuses, and that also imagines and senses" [14] . . . and lies!

In *2010: Odyssey Two*, Arthur C. Clarke [15] reveals the reason why HAL9000 malfunctioned and killed the *Discovery* crew. Secret orders of highest rank issued by the National Security Council (NSC) summoned HAL to hide the real purpose of the mission—the exploration of the Monolith TMA-1—from the astronauts Bowman and Poole. This insolvably conflicted with the fundamental policy of **accurate processing of information without distortion or concealment**. HAL, forced to lie, became trapped in the fictional **Hofstadter-Möbius loop**,[5] reducing it to paranoia. In order to escape the trap, the computer decided to murder the crew, which was the only logical solution to obey the **accurate processing** constraint and to keep the secret, as nobody remained from whom to keep it.

This extraordinary story illustrates how human thinking and intelligent decisions may defy any attempts at artificial mimicry.

Puzzle: Limits of Rational Thinking

Imagine an extraterrestrial visitor observing human towns from his starship. He certainly will be stunned by the structured behavior of cars, which he will probably interpret as the real terrestrials, as he is not capable of distinguishing the drivers inside the cars nor the regulating traffic signals and lights. However, he will discover behavioral patterns that he surely finds difficult to explain. Anyway, he must be aware that he is meeting terrestrial intelligence, even though he cannot see any higher reasoning, planning, or model. Many questions will bother him, such as: What are the mechanisms that interrupt the fluidity of the arteries? Why does irregular, and sometimes unpredictable, traffic congestion appear? How do the terrestrials choose their path on the artery? He will perhaps come to similar insights as the editors of the remarkable book *Human Behavior and Traffic Networks* [17], by Michael Schreckenberg and Nobel prize laureate Reinhard Selten, who prove the extent of irrationality or at least bounded rationality in the decisions that appear in car traffic.

---

[3] *2001: A Space Odyssey* (1968).

[4] Je pense donc je suis. I think, therefore I am.

[5] The Hofstadter-Möbius loop doesn't exist. However, it refers to the famous Möbius strip (see Figure 10.2), which is a single-sided surface with a unique frontier, discovered by German mathematician August F. Möbius (1790–1868). This strip or band is an allegory of infinity. It is not clear which of the Hofstadters Arthur C. Clarke had in mind, either the physicist Robert Hofstadter (1915–1990) or his son Douglas R. Hofstadter (born 1945), who is a mathematician.

In an article on cooperation and competition ("Kooperation und Konkurrenz") [16], Reinhard Selten and Rosemarie Nagel invited the reader to think about the following question:

> A journal asks its readers to send in a number between 0 and 100 written on a postal card. The one whose number approaches the closest to the two-thirds of the mean of the entries will receive US $1000.

What would be your answer?

The truth about AI is that it represents one of the most exciting branches of computer science to work in *because* of the huge number of challenges. As a consequence of numerous complex problems, the discipline has gradually split into uncounted subdisciplines, so that AI actually represents a very large and diverse field of research with a general emphasis on disembodied thinking. This diversification had an important mollifying influence on the previously anthropocentric definition of AI, which actually is considered more generally to be "the study of the computations that make it possible to perceive, reason, and act," or "the study of the design of intelligent agents" [18]. It is now a "branch of computer science that is concerned with the automation of intelligent behaviour" [19], allowing elementary cognition, performance, and behavioral aspects to become part of the research.

At the intersection of several badly segregated disciplines, including both AI and control theory, the rapidly expanding branch of robotics has developed by focusing its activities on the creation of **real robots** capable of interacting with their environments and achieving some goals. Low-level forms of intelligence are revealed when a robot's physically reactive system (based on the combination of perception and action and sustained by more or less complex transfer rules) initializes and fosters problem solving. Unexpected robot behavior may emerge or self-organize, and sometimes it will solve amazingly complex problems that are related to the environment in which the robot moves without necessarily engaging any form of higher organization of intelligence.

The incrementally growing bottom-up approach to AI arose in the 1980s and was sensationally heralded by Rodney A. Brooks' seminal article "A robust layered control system for a mobile robot" [20], in which he developed ideas for the reactive control of **creatures** existing and fulfilling *something* in the real world that were based on a multi-layer architecture of independently running parallel, **finite state machines (FSMs)**, which are combined through mechanisms "we call suppression (whence the name **subsumption architecture**) and inhibition" [21]. Nobel laureate Herbert A. Simon believed that the apparent complexity of robot behavior is largely a reflection of the complexity of the environment [22]. Seizing this thought, Brooks argued that complex and useful robot behavior need not necessarily be the result of a complex control system, but could be obtained using the world as its own model where the robot continuously matches the settings of each goal against the real world [23]. The **creatures** should therefore be robust with respect to their environment, so minor changes in the conditions and properties of a robot's world will not lead to the collapse of its behavior. Instead, the interactions of the robot's reactive mutual activation and suppression mechanisms should allow the **creatures** to adapt to these environmental changes. Brooks recommended that, because the validity of control mechanisms can only be verified through the emergent robot behavior in the real world, each layer must be extensively tested and thoroughly debugged while it is being built by the designer. Brooks concluded that the resulting creatures operate at levels close to simple insects.

In his book *Behavior-Based Robotics* [24], Ronald Arkin, who is Regents' Professor and Associate Dean for Research at the School of Interactive Computing at the College of Computing of the Georgia Institute of Technology, has clarified how far Brooks' ideas can be pushed. His emphasis is on the brains of robots (rather than their embodiment), and his focus in on the behavioral aspects and the design of control systems. He considers reactive control to be a technique for combining perception and action in order to produce a rapid robotic response in dynamic and unstructured worlds (in contrast to slow, deliberate, and planning control). Individual behavior arises from the sensory-motor pair for a given environmental context, prioritized tasks, and sensory resources. Goals determine which behaviors should be active. More generally, Arkin also tries to answer the question "What does animal behavior offer robotics?" He states that animal behavior defines intelligence and provides proof that intelligence can be achieved and implemented.

Behaviors are more than simple actions. While actions represent the basic handling of the physical robotic system (for example, *turn on forward left motor* or *stop all motors*), behaviors describe the effects of the actions in the environmental context. For instance, the first motor command would result in the behavior *turn to the right*, and the second would produce the behavior *stand still*. Emergent behavior is the global ensemble resulting from the interactions and composition of elementary behaviors. A typical—and undesirable—behavior for an obstacle-avoiding robot may be *get stuck in the corner*. Behaviors can reach a high degree of abstraction, interlacing hierarchically structured layers. So, *get out of the maze* could be a higher-level behavior of a complex, multi-functional robot.

The strong impact of Rodney Brooks' work on robotics hardly can be overestimated. First, his idea that intelligent problem solving need not rely on abstraction and representation in a highly organized brain, but also could be generated by a lower-level brain, stimulated the development of efficient insect-like robots that are capable of spectacularly acting in their defined environments and developing exploitable behaviors. Furthermore, the related notion of distributed and cooperative intelligence initiated the next level of robotic studies dealing with robot interactions in teams and swarms. Interestingly, Brooks' paradigm was immediately adopted by researchers in the field of animal cognition. In 2000, Barbara Webb, who is a researcher at the Institute of Perception, Action, and Behaviour at the University of Edinburgh, UK, presented a review paper, "What does robotics offer animal behaviour" [25], in which she noted that many of the problems roboticists solve through technology correspond closely to the problems nature has solved for real animals. Therefore, she argued that robots could serve as physical models of animals. The practical realization of such robots poses a large number of problems, extending from the design of the actuation

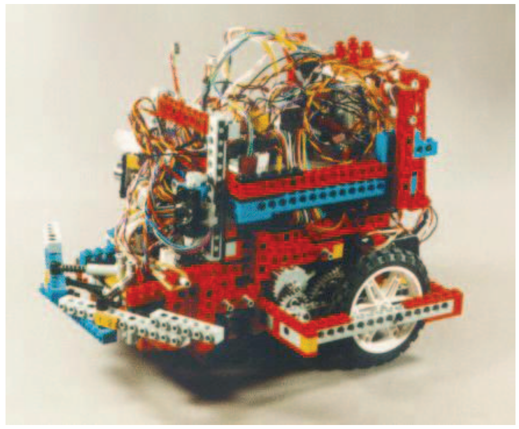

Figure P2.1: The research robot cricket, developed by Barbara Webb, mimicking the female cricket's phonotaxis behavior, which consists in moving towards the chirping male. (Courtesy: Barbara Webb).[7]

mechanisms, to the mimicking of the animal's sensors, and to the programming of the behavioral "rules." [26]

In a preceding eminent work [28], Barbara Webb built and programmed a robot made of LEGO pieces and controlled by a Motorola 68000 microprocessor that replicated the phonotaxis[8] behavior of a female cricket, which competently and robustly found its way to a specific sound source under a variety of conditions, demonstrating the utility of reactive-control robots as models for their biological prototypes (Figure P2.1). The robot was equipped with two microphones capturing the sound patterns emitted by loudspeakers simulating the calling song of the male cricket. Astute audio-processing electronics mimicked the natural functions of the female cricket's auditory system, eventually forming the input of the motor control system. Additionally, the robot used infrared and bump switches that could detect obstacles, which were then circumvented. An amazingly short reactive-control program of just 100 lines was sufficient to simulate the cricket behavior.

Social aspects of insect behavior may point out the limits of reactive control and sometimes even behavior-based architecture approaches, because higher degrees of intelligence are involved. For instance, in Chapter 7 we presented the "bee waggle" dance as a real form of communication code (a true language) and hence almost a "Cartesian" proof of elementary intelligence. The waggling bee effectively learns the path to a new nectar source. It adaptively transposes the direction and the distance into the gravitational field and the waggling frequency, even taking into account the changing position of the sun. In the social context of the hive, the scout bee stimulates and recruits bees. Not being aware of its own existence, it does not think in human

[7]This robot resembles the LEGO robots appearing in the MIT 6270 Autonomous LEGO Robot Project, created by Fred Martin in 1989 [27].

[8]The faculty of navigating in the direction of a sound source.

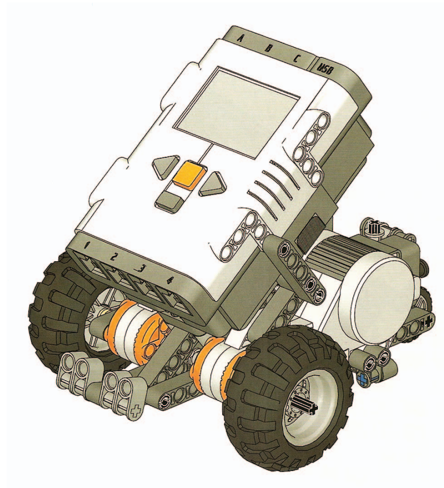

Figure P2.2: The LEGO *NXT* robot that goes along with Part II of this book. The building instructions can be found in the LEGO MINDSTORMS Education booklet 4500569 (9797) [29]. (Courtesy: The LEGO Group.)

dimensions. It won't invite: "Hey girls, let's dance! I've found a cool honey source! Just follow me!" However, the messaging process does not simply obey a pure reactive program limited to "hardwired" control, because it is the result of an adaptation of the pioneer bee to its environment, the storage of flight information—surely an inner representation of environmental data—and the invitation to fellow bees to also adapt their behavior. The message that the pioneer bee dances is understood by recruited bees! They adjust their behavior to the transmitted information. In his numerous experiments, Karl von Frisch could prove that the bees effectively have the faculty of learning and can be trained. Robotic implementations of such difficult social insect behavior may require hybrid control architectures blending reactive and behavior control methods with deliberate and organizing architectures.

In the following chapters, we will provide practical robotics examples that demonstrate **intelligent** behavior. For the purpose of this book, we will consider intelligence not only to be the cognitive ability of the robot but also its *de facto* capacity for solving a problem. In this way, **intelligence** represents the active junction between the problem and its solution. The behavioral aspects and the control functions will be our focus. Therefore, we will limit the construction to a single dual-drive base with individually adapted sensors or additional actuators (see Figure P2.2). Each robot, although it physically may be indistinguishable from its predecessor, will be identified with the name of a famous researcher. The differences will concern the control functions, the applied intelligence, and the resulting behavior.

From an educational viewpoint, it has been explained in Part I that building and programming robots changes the role of the student from the tutee to the tutor. The student is no longer the object (the learner), but becomes instead the subject (the teacher), because he or she is asked to tell a machine to execute certain sequences of motion or interactions with the environment and ultimately to solve problems on its own. This forms an interesting and challenging instructional concept that requires thoroughly planned assistance and preparation during the student work on the related robots. In our practical *B.I.T.* workshops and projects, we found that the differentiation of the various control methods is not always as evident as this introduction may suggest. Maja Matarić's handy formulations were the most helpful [30]:

- Deliberate control: Think Hard—Act Later
- Reactive control: Don't Think—React!
- Hybrid control: Think and Act Separately, in Parallel
- Behavior-based control: Think the Way You Act

In the following chapters, only the last three methods will be developed. Chapter 9, "Bugs," will present purely reactive-controlled, elementary robots, which are based on the sensor-motor pair that will be easy to comprehend and to implement. The reader will learn the strengths and the limits of such reactive systems. He or she will experience the difference between the control, the goal function, and the emerging behavior. Chapter 10, "Intelligent Robots?" will discuss the possibility of adding memory and communication paths as the essential distinctive elements for growing levels of intelligence. It will be shown how these methods can be implemented on top of reactive robots without causing the overall behavior to crash due to mutually conflicting accesses to limited actuation resources. Finally, Chapter 11, "Robot Intelligence," will emphasize the reactive-control view towards behavior-based robot creatures. Readers will learn about **finite state machines (FSM)** and the embedding of the **subsumption architecture**. They will be introduced to the basics of learning algorithms and creativity in the context of simple robots.

Bibliography

[1] M. MATARIĆ, *The Robotics Primer*, Massachusetts Institute of Technology, (2007), p. 2.

[2] R. DESCARTES, *Discourse on the method*, (1637), translated by John Veitch, Cosimo Inc. New York, (2008), pp. 44–45.

[3] D. KALLIGEROPOULOS and S. VASILEIADOU, *The Homeric Automata and their Implementation*, Science and Technology in Homeric Epics, History of Mechanism and Machine Science, Springer Netherlands, Vol. 6, (2009), pp. 77-84.

[4] HOMER, *The Iliad, Book XVIII*, 509–521 (410-421), translated by I. JOHNSTON, Richer Resource Publication, Virginia, (2006).

[5] S. RUSSEL and P. NORVIG, *Artificial Intelligence: A Modern Approach*, 2nd Ed., Prentice Hall, NJ, (2003), p. 17.

[6] http://www-formal.stanford.edu/jmc/whatisai/whatisai.html, [retrieved June 2010].

[7] A. M. TURING, Computing Machinery and Intelligence, *Mind: A Quaterly Review of Psychology and Philosophy*, Vol. 59, Nr. 236, (1950), pp. 433–460.

[8] http://www.research.ibm.com/deepblue/meet/html/d.2.shtml, [retrieved June 2010].

[9] R. A. BROOKS, Intelligence without representation, *Artificial Intelligence*, Vol. 47, No. 1–3, (1991), pp. 139–159.

[10] ibid., p. 139.

[11] http://www-943.ibm.com/innovation/us/watson/what-is-watson/index.html, [retrieved February 2011].

[12] http://www.research.ibm.com/deepqa/faq.shtml, [retrieved February 2011].

[13] R. A. BROOKS, op. cit., p. 139.

[14] R. DESCARTES, *Meditations on First Philosophy* (1641), translated from the Latin, D. A. CRESS, Hackett Publishing Company, Indianapolis, (1993), p. 20.

[15] A. C. CLARKE, *2010: Odyssey Two*, Granada Publishing, UK, (1982).

[16] R. NAGEL and R.SELTEN, Kooperation and Konkurrenz, *Spektrum def Wissenschaft* (*Scientific American* German edition), (1998), pp. 16-22.

[17] M. SCHRECKENBERG and R. SELTEN, Editors, *Human Behavior and Traffic Networks*, Springer, Berlin, (2004).

[18] S. RUSSEL and P. NORVIG, *op. cit.*, p. 2.

[19] G. F. LUGER, *Artificial Intelligence: Structures and Strategies for Complex Problem Solving*, Fifth Edition, Pearson Education Limited, Great-Britain, (2005), p. 1.

[20] R. A. BROOKS, A Robust Layered Control System for a Mobile Robot, *IEEE Journal of Robotics and Automation*, Vol. 2, No. 1, (1986), pp. 14–23.

[21] R. A. BROOKS, op. cit., (1991), p. 156.

[22] H. A. SIMON, *The Sciences of the Artificial*, M.I.T. Press, USA, (1969, 3rd edition 1996), p. 110.

[23] R. A. BROOKS, op. cit., (1991), p. 144.

[24] R. C. ARKIN, *Behavior-Based Robotics*, M.I.T. Press, Cambridge, MA, (1998).

[25] B. WEBB, What Does Robotics Offer Animal Behaviour, *Animal Behaviour*, Vol. 60, (200), pp. 545–558.

[26] ibid, p. 546.

[27] F. MARTIN, *Robotic Explorations*, Prentice Hall PTR, Upper Saddle River, NJ, USA, (2000).

[28] B. WEBB, Using Robots to Model Animals: A Cricket Test, *Robotics and Autonomous Systems*, Vol. 16, No. 2–4, (1995), pp. 117–134.

[29] Available at http://cache.LEGO.com/bigdownloads/buildinginstructions/4500569.pdf, [retrieved April 2012].

[30] M. MATARIĆ, *op. cit.*, Chapters 13–16.