1

Robotics and Artificial Intelligence

The Present and Future Visions

Sami Haddadin and Dennis Knobbe

INTRODUCTION

The rise of artificial intelligence is mainly associated with software-based robotic systems such as mobile robots, unmanned aerial vehicles, and increasingly, semi-autonomous cars. However, the large gap between the algorithmic and physical worlds leaves existing systems still far from the vision of intelligent and human-friendly robots capable of interacting with and manipulating our human-centered world. The emerging discipline of machine intelligence (MI), unifying robotics and artificial intelligence, aims for trustworthy, embodiment-aware artificial intelligence that is conscious both of itself and its surroundings, adapting its systems to the interactive body it is controlling. The integration of AI and robotics with control, perception and machine-learning systems is crucial if these truly autonomous intelligent systems are to become a reality in our daily lives. Following a review of the history of machine intelligence dating back to its origins in the twelfth century, this chapter discusses the current state of robotics and AI, reviews key systems and modern research directions, outlines remaining challenges and envisages a future of man and machine that is yet to be built.

1.1 MACHINE INTELLIGENCE: HISTORY IN A NUTSHELL

1.1.1 *Back to the Roots*

The basic vision of robotics and AI can be traced back to twelfth-century Europe.[1] Literature from this period mentions a mystical creature called the golem, which had a human-like shape but was significantly stronger than a normal human. The

[1] Wöll, "Der Golem: Kommt der erste künstliche Mensch und Roboter aus Prag?" in Nekula, Koschmal, and Rogall (eds), *Deutsche und Tschechen: Geschichte - Kultur - Politik* (Beck 2001) 233–245.

golem was described as a harmless creature used by its creator as a servant. In the legend of the golem of Prague, first written down at the beginning of the nineteenth century, Rabbi Löw created the golem to relieve him of heavy physical work and to serve humans in general.[2] The real-world realization of this idea had a long way to go.Some of the earliest scientific writings relating to machine intelligence date back to the fifteenth century, the period of the Renaissance. Leonardo da Vinci (1452–1519), the universal savant of his time,[3] decisively influenced both art and science with a variety of inventions, including, for example, a mechanical jumper, hydraulic pumps, musical instruments, and many more. However, the two inventions that stand out from a robotics point of view were Leonardo's autonomous flying machine and his mechanical knight, also known as Leonardo's robot.[4] The latter is a mechanism integrated into a knight's armor, which could be operated via rope pulls and deflection pulleys, enabling it to perform various human-like movements – clearly first steps in robotics. Wilhelm Schickard (1592–1635)[5] developed and built the first known working mechanical calculator. It was a gear-based multiplication machine that was also used for some of Kepler's lunar orbit calculations.

Sir Isaac Newton (1642–1726), one of the world's greatest physicists, is best known for laying the foundations of classical physics by formulating the three laws of motion.[6] He was also an outstanding mathematician, astronomer and theologian. In the field of mathematics, he developed a widely used technique for solving optimization problems (nowadays called Newton's method) and founded the field of infinitesimal calculus. Gottfried Wilhelm Leibniz (1646–1716) worked in parallel with Newton on this topic but conceived the ideas of differential and integral calculus independently of Newton.[7] Leibniz, who is known for various other contributions to science, is often referred to as one of the first computer scientists due to his research on the binary number system. Slightly later, Pierre Jaquet-Droz (1721–1790) built amazing mechanical inventions such as The Writer, The Musician and The Draughtsman.[8] The Draughtsman, for example, is a mechanical doll that draws with a quill pen and real ink on paper. The input device was a cam disk that essentially functions as a programmable memory defining the picture to be drawn. With three different cam disks, the The Draughtsman was able to draw four different artworks. In addition to these fascinating machines, Jaquet-Droz and his

2 Grün and Müller, *Der hohe Rabbi Löw und sein Sagenkreis* (Verlag von Jakob B Brandeis 1885).
3 Grewenig and Otto, *Leonardo da Vinci: Künstler, Erfinder, Wissenschaftler* (Historisches Museum der Pfalz 1995).
4 Moran, "The da Vinci Robot" (2006) 20(12) *Journal of Endourology* 986–990.
5 Nilsson, *The Quest for Artificial Intelligence* (Cambridge University Press 2009).
6 Westfall, *Never at Rest. A Biography of Isaac Newton* (Cambridge University Press 1984).
7 Nilsson (n 5).
8 Soriano, Battaïni, and Bordeau, *Mechanische Spielfiguren aus vergangenen Zeiten* (Sauret 1985).

Cambridge University Press
978-1-108-42482-0 — Algorithms and Law
Edited by Martin Ebers , Susana Navas
Excerpt
More Information

business partner Jean-Frédéric Leschot later started to build prosthetic limbs for amputees.

Another memorable figure in the history of machine intelligence is Augusta Ada Byron King (1815–1852).[9] The Countess of Lovelace is known to be one of the first to recognize the full potential of a computing machine. She wrote the first computer program in history, which was designed to be used for the theoretical analytical engine proposed by Charles Baggage. The programming language Ada was named after her. These fundamental technological advances in the areas of mechanics, electronics, communications and computation paved the way for the introduction of the first usable computing machines and control systems, which began around 1868. The first automatic motion machines were systematically analyzed, documented, reconstructed, and taught via collections of mechanisms.

A mechanism can be defined as an automaton that transforms continuous, typically linear, movements into complex spatial motions. Ludwig Burmester (1840–1927) was a mathematician, engineer and inventor, and the first person to develop a theory for the analysis and synthesis of motion machines.[10] Later in this period, Czech writer and dramatist Karel Čapek (1890–1938) first used the word "robot" in his science-fiction work. The word "robot" is derived from *robota*, which originally meant serfdom, but is now used in Czech for "hard work." Through his 1920 play *R.U.R.* (*Rossums Universal Robots*), Čapek spread his definition of robot to a wider audience.[11] In this play, the robots were manufactured to industry standards from synthetic organic materials and used as workers in industry to relieve people from heavy and hard work.

We now come to the pre-eminent philosopher and mathematician Norbert Wiener (1894–1964). From his original research field of stochastic and mathematical noise processes, he and his colleagues Arturo Rosenblueth, Julian Bigelow and others founded the discipline of cybernetics in the 1940s.[12] Cybernetics combines the analysis of self-regulatory processes with information theory to produce new concepts, which can be said to be the precursors of modern control engineering, thus building significant aspects of the theoretical foundations of robotics and AI. Wiener developed a new and deeper understanding of the notion of feedback, which has significantly influenced a broad spectrum of natural science disciplines. Alan Turing (1912–1954) worked in parallel with Wiener in the field of theoretical computer science and artificial intelligence.[13] Most people interested in artificial intelligence today are familiar with his name through the Turing test. This test was

9 Nilsson (n 5).
10 Koetsier, "Ludwig Burmester (1840–1927)" in Ceccarelli (ed), *Distinguished Figures in Mechanism and Machine Science*, *History of Mechanism and Machine Science*, vol 7 (Springer 2009) 43–64.
11 Nilsson (n 5).
12 Ibid.
13 Ibid.

devised to determine whether a computer or, more generally a machine, could think like a human. His groundbreaking mathematical model of an automatic calculating machine that can solve complex calculations is today known as a Turing machine. The Turing machine models the process of calculating in such a way that its mode of operation can be easily analyzed mathematically, making the terms "algorithm" and "computability" mathematically manageable for the first time.

A similarly renowned researcher and colleague of Turing was John von Neumann (1903–1957).[14] He developed the von Neumann computer architecture, which still forms the basis of the operation of most computers today. As well as collaborating with Turing on AI research, he also worked on other mathematical topics like linear programming and sorting programs. Von Neumann's concept of self-reproducing machines, developed in 1940, testifies to his outstanding capabilities.[15] The aim of this concept was to describe an abstract machine, which, when in operation, replicates itself. To achieve this goal von Neumann also developed the concept of cellular automata. According to von Neumann, a cellular automaton is a collection of states in a two-dimensional grid of cells, which forms a certain pattern. A cell represents one of twenty-nine possible states, which can change over time. The change of state of a cell is determined by the states of the neighboring cells from the previous time step as input. The theory of cellular automata defined the elementary building blocks responsible for the concept of self-replicating machines. With these building blocks, von Neumann created the universal constructor, which is a particular pattern of different cell states. This pattern contains three different sub-units: an information carrier for storing its own construction plan, a construction arm, which builds itself up in the free grid according to the construction plan, and a copying machine for copying the construction plan. This made it possible for von Neumann to develop a self-replicating machine within the concept of cellular automata.

A famous mathematician and inventor who also worked in the field of digital computing is Claude Elwood Shannon (1916–2001). His groundbreaking ideas on logical circuit design for digital computers and information theory had an enormous impact on the research community of his time, and continue to do so today. In 1948, with his book *A Mathematical Theory of Communication*,[16] he laid important foundations for today's high-speed telecommunications and data processing by mathematically tackling the problem of data transmission via a lossy communication channel. He developed a coding algorithm that made it possible to restore the originally transmitted information from previously coded lossy data. In a further

[14] Ibid.
[15] Von Neumann and Burks, "Theory of Self-Reproducing Automata" (1966) 5(1) *IEEE Transactions on Neural Networks* 3.
[16] Shannon, "A Mathematical Theory of Communication" (1948) 27(3) *Bell System Technical Journal* 379–423.

publication,[17] he developed a complete theory of channel capacity, which defined the maximum data rate that can be transmitted lossless over a specific communication channel type. In 1949, he published the formal basics of cryptography, thus establishing it as a scientific discipline.[18]

At the beginning of 1941, the engineer and computer scientist Konrad Zuse (1910–1995) made headlines with the world's first functional programmable digital computer, the Z3, built in cooperation with Helmut Schreyer.[19] Zuse also demonstrated that machines can assemble themselves on a variable scale, long before the idea of robotic assembly systems had been conceived.[20] Based on John von Neumann's ideas and proofs that it is theoretically possible to build a machine that can reproduce itself, Zuse published his implementation ideas for such a machine in the journal *Unternehmensforschung* under the title "Gedanken zur Automation und zum Problem der technischen Keimzelle" ("Thoughts on Automation and the Problem of the Technical Germ Cell").[21] In the 1970s he designed the assembly robot SRS72 in his own construction workshop as a functional demonstration of this idea. The SRS72 machine could automatically assemble prefabricated manually supplied parts by positioning two work pieces and connecting them with screws. This prototype machine was the starting point for a complete self-reproducing system. According to Zuse, an entire automated workshop is required to perform all the complex manufacturing and assembly steps necessary to obtain a self-producing system.[22]

Independently of Zuse, the physicist Richard Phillips Feynman (1918–1988) also studied von Neumann's ideas. His own research area was quantum field theory, and he was awarded the Nobel Prize in 1965 for his work on quantum electro dynamics. Today, however, he is also regarded as a visionary of self-reproducing machine technology. His famous lecture, "There's Plenty of Room at the Bottom," on the future opportunities for designing miniaturized machines that could build smaller reproductions of themselves was delivered in 1959 at the annual meeting of the American Institute of Physics at the California Institute of Technology and published the following year in the journal *Engineering and Science*.[23] Feynman's speech is frequently referenced in today's technical literature in the fields of

[17] Shannon, "Communication in the Presence of Noise" (1949) 86 *Proceedings of the IRE* 10–21. 10.1109/JRPROC.

[18] Shannon, "Communication Theory of Secrecy Systems" (1949) 28(4) *Bell System Technical Journal* 656–715.

[19] Bauer et al., *Die Rechenmaschinen von Konrad Zuse* (Springer 2013).

[20] Eibisch, "Eine Maschine baut eine Maschine baut eine Maschine…" (2011) 1 *Kultur und Technik* 48–51.

[21] Zuse, "Gedanken zur Automation und zum Problem der technischen Keimzelle" (1956) 1(1) *Unternehmensforschung* 160–165.

[22] Ibid.

[23] Feynman, "There's Plenty of Room at the Bottom," talk given on 29 December 1959 (1960) 23 (22) *Science and Engineering* 1–13.

micro- and nanotechnology, which speaks for the high regard in which his early vision is held in expert circles.

Very few people had the knowledge and skills to program complex early computing machines like the Z3 computer. Unlike today's programming languages that use digital sequence code, these machines were programmed with the help of strip-shaped data carriers made of paper, plastic or a metal-plastic laminate, which store the information or the code lines in the punched hole patterns. One person who mastered and shaped this type of programming was American computer scientist Grace Hopper (1906–1992).[24] She did not work with the Z3, but on the Mark I and II computers she designed the first compiler called A-0. A compiler is a program that translates human readable programming code into machine-readable code. She also invented the first machine-independent programming language, which led to high-level languages as we know them today.

Returning to robotics in literature, a short story that still exerts a powerful influence on real-world implementation of modern robotics and AI systems as we know them today is Isaac Asimov's (1920–1992) science-fiction story "Runaround," published in 1942, which contained his famous "Three Laws of Robotics":[25]

> One, a robot may not injure a human being, or, through inaction, allow a human being to come to harm. [. . .] Two, a robot must obey the orders given it by human beings except where such orders would conflict with the First Law. [. . .] And three, a robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.

Asimov's early ideas, including his vision of human–robot coexistence, paved the way for the concept of safety in robotics. Asimov's Three Laws, formulated as basic guidance for limiting the behavior of autonomous robots in human environments, are enshrined, for example, in the Principles of Robotics of the UK's Engineering and Physical Sciences Research Council (EPSRC)/Art and Humanities Research Council (AHRC), published in 2011.[26] These principles lay down five ethical doctrines for developers, designers and end users of robots, together with seven high-level statements for real-world applications.

Shortly before the vast technological advancements in the second half of the twentieth century began, the first rudimentary telerobotic system was developed in 1945 by Raymond Goertz at the Argonne National Laboratory.[27] It was designed to control, from a shelter, a robot that could safely handle radioactive material. From the 1950s on, the first complex electronics were developed, further optimized and miniaturized, and modern concepts of mechanics were created. The first

[24] Beyer, *Grace Hopper and the Invention of the Information Age* (BookBaby 2015).
[25] Asimov, *Astounding Science Fiction*, chapter "Runaround" (Street & Smith 1942).
[26] Prescott and Szollosy, "Ethical Principles of Robotics" (2017) 29(2) *Connection Science* 119–123.
[27] Goertz and Thompson, "Electronically Controlled Manipulator" (1954) 12 *Nucleonics (US)* 46–47.

mechatronic machines, such as fully automated electric washing machines[28] or the first industrial robots,[29] were invented, and the concept of AI was further developed. Through the mathematical work of Jacques S Denavit (1930–2013), Richard Hartenberg (1907–1997) and Rudolf August Beyer (1892–1960), one of the most important methods of calculating the direct kinematics of robots was developed around the year 1955.[30] This matrix calculus, known today as the Denavit–Hartenberg Convention, calculates how the joints of a robot have to be adjusted in order for it to be able to approach a specific point in space. In the same year, John McCarthy (1927–2011), an American cognitive computer scientist and inventor of the famous programming language Lisp, introduced the term "artificial intelligence."[31] He also organized the famous Dartmouth Conference in the summer of 1956, which is considered the birth of AI as a research field.

Marvin Lee Minsky (1927–2016) was an American mathematician and cognitive scientist as well as a colleague of McCarthy in the same AI working group at Massachusetts Institute of Technology (MIT).[32] He is known for the invention of head-mounted graphical displays and for his work in artificial neural networks. Together with Seymour Papert, he wrote the book *Perceptrons*, which is still required reading for the analysis of artificial neural networks. He introduced several famous AI models and developed SNARC, the first neural network simulator. The late 1950s can also be seen as an important opening stage in the modern theory of optimization and optimal control. The field of optimal control deals with the process of calculating appropriate control laws for a given system in order to meet certain desired optimality criteria. In this context, at the end of the 1950s the mathematicians Lev Semyonovich Pontryagin (1908–1988) and Richard E Bellman (1920–1984) published a series of new fundamental optimization methods, such as Pontryagin's maximum principle,[33] Bang-Bang control,[34] the Hamilton–Jacobi–Bellman equation or the Bellman equation for dynamic programming,[35] which changed the entire field of mathematical optimization and control. These advances continue to this day to have a major influence on various practical areas from engineering to economics.

28 Milecki, "45 Years of Mechatronics–History and Future" in Szewczyk, Zieliński, and Kaliczyńska (eds), Progress in Automation, Robotics and Measuring Techniques in Szewczyk, Zieliński, and Kaliczyńska (eds), Progress in Automation, Robotics and Measuring Techniques (Springer 2015).
29 Nilsson (n 5).
30 Denavit and Hartenberg, "A Kinematic Notation for Lower-Pair Mechanisms Based on Matrices" Trans. of the ASME (1955) 22 *Journal of Applied Mechanics* 215–221.
31 Nilsson (n 5).
32 Ibid.
33 Boltyanskii, Gamkrelidze, and Pontryagin, "Towards a Theory of Optimal Processes" (in Russian) (1956) 110(1) *Reports Acad Sci USSR* 1–10.
34 Pontryagin et al., *Mathematical Theory of Optimal Processes* (in Russian) 1961.
35 Bellman, *Dynamic Programming*, vol 295 (Rand Corp Santa Monica CA 1956); Bellman, *Dynamic Programming* (Princeton University Press 1957).

In 1957 the first autonomous underwater vehicle, the Self-Propelled Underwater Research Vehicle (SPURV), was invented at the Applied Physics Laboratory at the University of Washington by Stan Murphy, Bob Van Wagennen, Wayne Nodland, and Terry Ewart;[36] this system was used to measure the physical properties of the sea.

A few years later, in 1960, electrical engineer and mathematician Rudolf Emil Kalman (1930–2016) developed the Kalman filter in cooperation with Richard S Bucy and Ruslan L Stratonovich.[37] This mathematical algorithm is capable of predicting system behavior based on a dynamic model and suppressing additive noise at the same time. In the context of this algorithm Kalman introduced two new system analysis concepts: system observability and controllability.[38] The concept of observability analyzes how well the internal states of a system can be calculated by measuring its output. Controllability measures how an input signal changes the internal states of a system. These system analysis methods are crucial for the design of a Kalman filter, but also provide very important system information for the design of stable control loops in robots, process machines or driver assistance systems in cars. The Kalman filter itself is still one of the most important signal-processing tools in modern robotics, but is also used in various other disciplines such as AI, navigation, communications and macroeconomics.

The basic theories of robotics continued to expand, with developments in hardware and control, such as electric motor and sensor systems. In 1961 Joseph Engelberger (1925–2015), an American entrepreneur, physicist and engineer known as the father of industrial robots, developed, together with his company, the first industrial robot, Unimate.[39] A few years later, in 1964, a machine-learning algorithm called support-vector machine (SVM) was invented by mathematicians Vladimir Naumovich Vapnik and Alexey Yakovlevich Chervonenkis (1938–2014).[40] The original SVM algorithm is a linear classifier for pattern recognition. In 1992 the original method was extended to a nonlinear classifier by applying the so-called kernel trick;[41] the algorithm's final stage of development, still used today, was reached in 1995.[42]

[36] Van Wagenen, Murphy, and Nodland, *An Unmanned Self-Propelled Research Vehicle for Use at Mid-Ocean Depths* (University of Washington 1963); Widditsch, *"SPURV-The First Decade" No APL-UW-7215*, Washington University Seattle Applied Physics Lab 1973.

[37] Kalman, "A New Approach to Linear Filtering and Prediction Problems" Transaction of the ASME (1960) 82(1) *Journal of Basic Engineering* 35–45.

[38] Kalman, "On the General Theory of Control Systems" (1960) *Proceedings First International Conference on Automatic Control*, Moscow, USSR.

[39] Nilsson (n 5).

[40] Chervonenkis, *Early History of Support Vector Machines. Empirical Inference* (Springer 2013); Vapnik and Chervonenkis, Об одном классе алгоритмов обученияраспознаванию образов (On a Class of Algorithms of Learning Pattern Recognition) (1964) 25(6) *Avtomatika i Telemekhanika*.

[41] Boser, Guyon, and Vapnik, "A Training Algorithm for Optimal Margin Classifiers" *Proceedings of the Fifth Annual Workshop on Computational Learning Theory* (ACM 1992) 144–152.

[42] Cortes and Vapnik, "Support-Vector Networks" (1995) 20(3) *Machine Learning* 273–297.

Back in 1966, the computer program ELIZA was developed and introduced at MIT's Artificial Intelligence Laboratory under the direction of Joseph Weizenbaum.[43] ELIZA is a program for natural language processing that uses pattern matching and substitution methodologies to demonstrate communication between humans and machines by simulating a coherent conversation. Three years later American engineer Victor Scheinman (1942–2016) designed the first successful electrically operated, computer-controlled manipulator.[44] This robotic arm had six degrees of freedom, and was light, multi-programmable and versatile in its motion capabilities. Later on, the robot was amended for industrial uses such as spot welding for the automotive industries. In the field of machine learning, David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams introduced the modern version of the backpropagation algorithm in 1968.[45] This method is used in artificial neural networks to train networks and is a standard tool in this field today.

1.1.2 *The Modern Era of Robotics and AI*

The modern era of robotics and AI is characterized by ever greater miniaturization of electronics and mechatronics and an enormous increase in computing power, developments that have led to more practical robotic systems. The first humanoid robot to mimic human motion, the WaBot 1, was introduced by a Japanese research team from Waseda University in 1973.[46] WaBot 1 had very basic capabilities to walk, grab objects and transport them from one place to another. In 1978 Unimation released a new and more versatile version of the Unimate, called the Programmable Universal Machine for Assembly (PUMA).[47] PUMA has become very popular in industry and academia and over time has become an archetype for anthropomorphic robots. It remains widely used today as a reference example and benchmark system in academic robotics books and publications worldwide.[48]

In the 1980s the modern field of reinforcement learning was founded by combining different approaches from various disciplines. The starting point was the idea of trial-and-error learning, which was derived from psychological studies on animal learning dating from the early eighteenth century.[49] Reinforcement is the expression

43 Nilsson (n 5).
44 Scheinman, "Design of a Computer Manipulator" Stanford AI Memo AIM-92, 1 June 1969.
45 Rumelhart, Hinton, and Williams, "Learning Representations by Back-Propagating Errors" (1986) 323 *Nature* 533–536.
46 Kato, "of WABOT 1" (1973) 2 *Biomechanism* 173–214.
47 Beecher, *Puma: Programmable Universal Machine for Assembly, Computer Vision and Sensor-Based Robots* (Springer 1979).
48 Corke, "Robot Arm Kinematics" in Corke (ed), *Robotics, Vision and Control* (Springer 2017); Çakan and Botsali, "Inverse Kinematics Analysis of a Puma Robot by using MSC Adams" The VIth International Conference Industrial Engineering and Environmental Protection 2016 193–228.
49 Woodworth, *Experimental Psychology* (Holt 1938), Department of Psychology Dartmouth College Hanover, New Hampshire 1937; Woodworth, "Experimental Psychology (Rev edn)" (1954) 18(5) *Journal of Consulting Psychology* 386–387.

of a certain behavior pattern in connection with an interaction of an animal with its environment. The animal receives different stimuli in temporal correlation with its behavior, causing certain behavior patterns to persist even after the stimuli have subsided. From the technical point of view, this process can be described as an optimization problem with some stochastic features in terms of incomplete knowledge of the whole system. A further development of the optimal control framework already mentioned can be used to describe and solve such a system. One of the first to implement this idea was Witten, with his adaptive optimal control approach.[50]

Another important aspect of the rise of the modern theory of reinforcement learning is temporal-difference (TD) learning, the origins of which lie in animal learning psychology. It can be seen as either a subclass or an extension of the general reinforcement learning idea. In contrast to the standard reinforcement approach, in TD learning the learner's behavior or strategy is adjusted not only after receiving a reward, but after each action before receiving it, based on an estimate of an expected reward with the help of a state value function. The algorithm is thus controlled by the difference between successive estimates. In 1959 Arthur Samuel implemented this approach for the first time in his checkers-playing program.[51]

In 1983, a further development of this reinforcement learning algorithm, the so-called actor–critic architecture, was applied to the control problem of pole balancing.[52] The year 1989 can be described as the year of full integration of optimal control methods with online learning. The time difference and optimal control methods were fully merged in this year with Chris Watkin's development of the Q-Learning algorithm.[53]

In addition to reinforcement learning, the 1980s saw seminal work in robot manipulator control. Early in the decade John J Craig and Marc Raibert published a new hybrid control technique for manipulators. Their system made it possible to simultaneously satisfy the position and force constraints of trajectories, enabling compliant motions of robot manipulators.[54] In the mid-1980s, Neville Hogan developed impedance control for physical interaction,[55] which was an important

50 Witten, "An Adaptive Optimal Controller for Discrete-Time Markov Environments" (1977) 34(4) *Information and Control* 286–295.

51 Samuel, "Some Studies in Machine Learning Using the Game of Checkers" (1959) 3(3) *IBM Journal of Research and Development* 210–229.

52 Barto, Sutton, and Anderson, "Neuronlike Adaptive Elements That Can Solve Difficult Learning Control Problems" (1983) 5 *IEEE Transactions on Systems, Man, and Cybernetics* 834–846.

53 Watkins, *Learning from Delayed Rewards* PhD Thesis, King's College 1989.

54 Raibert and Craig, "Hybrid Position/Force Control of Manipulators" (1981) 103(2) *Journal of Dynamic Systems, Measurement, and Control* 126–133.

55 Hogan, "Impedance Control: An Approach to Manipulation: Part I – Theory, Part II – Implementation, Part III – Applications" (1985) 107 *Journal of Dynamic Systems, Measurement and Control* 1–24.