INTRODUCTION

> The neurones concerned in the behavior of a single man probably exceed in number by a thousand-fold all the telephone lines in the world [...] Even if we knew the exact arrangement of each neuron in a man's brain it would take a model as large as St. Paul's Cathedral to make them visible to the naked eye (E.L. Thorndike, 1914; 1919).

> To make a model which would reproduce all the behavior of a rat would require a mechanism probably as large as the Capitol in Washington (C.L. Hull, 1935).

> A large building could not house a vacuum tube computer with as many relays as a man has in his head, and it would take Niagara Falls to supply the power and Niagara River to cool it. ENIAC with some 10000 tubes has no more relays than a flat worm (W.S. McCulloch, 1949).

The 'discovery' in the title of this book refers to the coming of age of a new concept of the machine. This concept is at the core of a methodology that profoundly influenced the sciences of the mind and behavior in the twentieth century. The main ambition of this methodology was to overcome, through that new concept of the machine, traditional oppositions between the inorganic and organic worlds, between the laws that govern the behavior of physical systems and those that govern the behavior of organisms, and between causal and teleological explanation.

The origins of this methodology are usually traced back to the middle of the 1940s, with the advent of cybernetics, which Norbert Wiener described, in his 1948 book, as the study of "control and communication in the animal and the machine." That is when the new machines, those equipped with automatic control and with forms, albeit primitive, of self-organization, seemed to suggest a way to escape those oppositions. Mechanical models were built to simulate adaptive behavior as well as various forms of learning of living organisms. The concepts of machine and teleology no longer seemed incompatible.

The fundamental insight of cybernetics, i.e. the proposal of a unified study of organisms and machines, was inherited, starting in the mid-1950s, by AI (Artificial Intelligence). However, AI proposed a different simulative methodology. To put it quite generally, this methodology used computer programs to reproduce performances which, if observed in human beings, would be regarded as intelligent. In the course of the quarter-century that followed, many cybernetics research programmes were cut back, if not altogether abandoned, in the face of the early successes of AI. This is recent history, and still more recent is the history of the revival of some cybernetic projects in the changed context of the 1980s. This revival brought about renewed interest in self-organizing systems, neural nets and connectionism, and influenced later developments in ALife (Artificial Life) and behavior-based and evolutionary robotics.

Some of the main points in this recent history are dealt with in Chapters 5 and 6 of the present book. One of my central claims is that certain basic features of the simulative methodology whose origins are usually put no further back than cybernetics, actually go back in significant ways to the early decades of the twentieth century. In the first four chapters of the book, while examining little-known, if not forgotten or unpublished, texts of the pre-cybernetic age, I investigate various projects involving artifacts that were considered to be mechanical models of organisms' behavior. I identify in those projects what Marvin Minsky called the "intellectual ancestors" of cybernetics, when he dated the latter's public emergence to the nearly simultaneous appearance, in 1943, of Rosenblueth, Wiener and Bigelow's paper on the teleological features of feedback machines, of McCulloch and Pitts' paper on neural nets, and of Craik's modeling approach in *The Nature of Explanation*. As a result of my investigation, the simulative approaches of the second half of the twentieth century should appear, in Chapters 5 and 6, in a light that should make it easier to grasp those elements, both of continuity and discontinuity, that they share with earlier stages in the discovery of the artificial.

In Chapter 1, taking my cue from the disagreement between Jacques Loeb and Herbert Jennings in the study of animal behavior at the beginning of the twentieth century, I bring to light some features of the discovery of the artificial. We shall see how, alongside the traditional opposition between reductionist and vitalist approaches in the study of organisms' behavior, another approach is adumbrated, in which the behavior of both "inorganic machines" and living organisms can be studied by referring to their functional organization. Furthermore, with the development of new technologies, it is suggested that a positive test for a behavioral theory is provided by building a machine that functions as the theory predicts. According to Loeb, the phototropic automaton (or "electric dog") built by John Hammond Jr. and Benjamin Miessner was precisely an "artificial machine" of this kind: it supported his own theory of phototropism in living organisms or "natural machines." As we shall see, Loeb's theory prompted the design or actual building of several phototropic robots: from Edward Tolman's "schematic sowbug" up to Walter Grey Walter's electronic "tortoises" and Valentino Braitenberg's "vehicles."

The heated disputes of Loeb's days about the opposition, in the study of animal behavior, between the world of intelligence and purpose on the one hand, and that of instincts and automatism on the other, now seem very distant. The same might be said of other more radical disputes, like that between vitalism and mechanism. And yet, various issues that were raised then with regard to both lower and higher animals were later raised with regard to machines. (After all, this is not by chance: had animals not been formerly considered as the true *automata*, lacking purposes and intentionality?) Think, for example, of issue of the legitimacy of describing animals in mentalistic or teleological terms or, if you like, of assuming towards them what Daniel Dennett calls the "intentional stance." As Jennings put it in 1906, "this [kind of description] is useful; it enables us practically to appreciate, foresee, and control [the] actions [of the animals] much more readily than we could otherwise do so."

Various issues recurring in those disputes are dealt with in Chapter 2. Early neurological and psychological associationism and connectionism (a term, it should be remembered, introduced by Edward Thorndike) is the framework for my analysis of early attempts to test hypotheses about the nature of memory, adaptation and learning through mechanical models. This kind of theory testing in the behavioral sciences appeared to the psychologist J.M. Stephens as a genuine *discovery*, when in 1929 he wrote about his "surprise" at finding the "startling possibilities" of building a working artifact embodying the hypotheses of his theory of learning. And the engineer S. Bent Russell, in discussing his mechanical model of simple forms of learning, had already pointed out that "as the cooperation of workers in different fields of knowledge is necessary in these days of specialists it may be argued that engineers can consistently join [with neurologists] in the consideration of a subject of such importance to man" as it is the study of nervous system. Bent Russell made this suggestion in 1913, about thirty years before the advent of cybernetics, and Wiener would have appreciated his words.

The "robot approach," promoted by Clark Hull in the context of behaviorism in the 1920s and 1930s, develops along the lines suggested by Bent Russell. The robot approach is the subject of Chapter 3. In this case, one is for the first time in the presence of a coherent research programme with an openly interdisciplinary calling, and of a fully aware simulative or modeling methodology, which raised the hope of progressively building 'intelligent' machines. Nicolas Rashevsky, who initially showed interest in Hull's approach, wrote in 1931 that a machine might be made that could not only express itself in a natural language and carry on a conversation but could also learn a language and even lie intentionally—even if the actual construction of such a machine would have taken "tremendous expense and labor."

In Chapter 4, I examine Kenneth Craik's position between Hull's robot approach and the new-born cybernetics, and discuss Craik's symbolic theory of thought. I also emphasize how various proposals for 'reconciling' natural science and teleology were discussed before the advent of cybernetics. These proposals appear to have been influenced by their author's acquaintance with some self-regulating artifacts, which were explicitly defined as *teleological* in the above-mentioned 1943 paper by Rosenblueth, Wiener and Bigelow.

I have already alluded to the content of Chapters 5 and 6, in which the notion of artificial as *computational* or *algorithmic*, not merely as *inorganic*, comes into focus with the advent of AI. Finally, in Chapter 7, while bringing out the main thread running through my entire investigation, I state some theses that I believe have characterized the different stages in the discovery of the artificial over the course of the twentieth century. I use these theses to identify the points of convergence and divergence between the different approaches and research trends discussed in my investigation, with the aim of understanding major critical points better. Furthermore, I try to identify in those theses the roots of certain conflicting positions that are still being debated today in much cognitive science and philosophy of mind, concerning, for example, functionalism, the role of representations, and intentional vocabulary.

The present book is not concerned with the history of those automata or robots that were built or imagined before the twentieth century. On the one hand, that history has already been told several times and from various points of view; on the other hand, the 'artificial' as understood in the present book is quite remote from the clockwork automata which were especially popular during the eighteenth century. Notwithstanding certain important intuitions of people like Jaques de Vaucanson, such automata could hardly be seen as pre-cybernetic machines endowed with actual sense organs, self-controlling or feedback devices and motor organs. They could not 'adjust' their responses to incoming stimuli. As for self-controlling automata, their early history was first carefully reconstructed by Otto Mayr. These machines are linked to the discovery of the artificial in so far as they exhibit some form of reactive behavior and, albeit in fairly elementary forms, some ability to change their internal organization when interacting with the environment on the basis of their previous 'history.' Hence, these machine exhibit different forms of adaptation, memory, and learning. However, machines of this kind, as well as other artifacts such as computer programs or contemporary mobile robots, are here chiefly analyzed as tools for theoretical investigations into intelligent behavior.

In fact, there are three stories, as it were, that are told in the present book. First, there is the story of building artifacts of different kinds and complexity whose behaviors, if exhibited by organisms, would be regarded as 'adapted,' 'learned,' 'intelligent,' and so forth. Second, there is the story of the attempts to build artifacts that could be considered as tools for building and testing theories about those behaviors. Third, there is the story of the gradually developing discussion about the implications of the first two stories for classical epistemological questions, the mind-body problem, the possibility of scientific theorizing about cognition, and the like. These three stories proceed in parallel, and I have tried to bring out salient conceptual connections between them.

Progress in electronics and computer science and technology allow current modeling approaches to obtain results that inevitably make pre-cybernetic machines, which play a leading role in much of the present book, seem rough and extravagant. These machines almost always are electrical circuits or electro-mechanical robots with fairly simple performances. Nonetheless, the wealth of interesting issues surrounding early stages of the discovery of the artificial is surprising. The present book shows how, with regard to different notions of the machine, many problems were raised which are still being dealt with by scientists and philosophers of mind currently working on organism's behavior and the human mind.

In fact, those pre-cybernetic machines, just like those of the cybernetic age and those in several current areas of research, were aimed at simulating *functions*, rather than at reproducing the *external appearances* of living organisms. The flight analogy argument ("if material is organized in a certain way, it will fly like an eagle; if it is organized in another way, it will fly like an airplane") was already stated by Hull, and would later recur in AI. Since many of those pre-cybernetic machines were built with the ambitious aim of testing hypotheses on organism behavior, they were designed to be working models embodying *theories*, whether psychological or neurological, about behavior. They neither surprise nor seduce us because of the fairly realistic appearance that characterized the automata of former centuries, which chiefly imitated the outward appearance of animals and human beings. On the contrary, their goal is an avowedly non-mimetic one. A behaviorist model of learning by conditioning does not have to salivate like Pavlov's dog; rather, it has to grasp some of the *essential features* of the learning phenomenon. As we shall see, this was the goal underlying the building of simulative models of behavior, a goal clearly stated from the outset: from the days of Hull and his early attempts to simulate conditioned reflex by means of electro-mechanical artifacts, continuing through the days of Grey Walter and his electronic tortoises, and up to much current AI and robotics.

Some of the pre-cybernetic machines described in the present book were rather carefully designed, as can be seen in the neat drawings by Bent Russell, or in the circuit diagrams by Norman Krim. Other machines are known to us through rough drawings, like those by Rashevsky, or through rather involved descriptions, like that of Hugh Bradner's mobile robot, or through informal descriptions, like that of Tolman's schematic sowbug. Still other machines were actually built, including those by Krim himself, by Hammond Jr. and Miessner, by H.D. Baernstein, and by Thomas Ross. Exhibited to the public, they were immediately dubbed "intelligent" or "thinking" machines.

The use of these mental terms in the case of machines that today seem so rudimentary stimulates reflections on their public impact and the exaggeration of their actual intelligent abilities—a phenomenon that occurred again and again until more recent times. A "thinking machine" was how William Ross Ashby's homeostat was called in the 1940s, and the very expression used for the first digital computers—"electronic brains"—speaks volumes in this regard. We can well imagine that, when Baernstein's electrical model of conditioned reflex was exhibited at the 1929 Convention of the Midwest Psychological Association, the impression it created was no less than the one created when Grey Walter's electronic tortoises were exhibited at the 1951 Festival of Britain, or when Arthur Samuel's computer program, which played successfully a game of checkers, was exhibited on the 1960 CBS television show *The Thinking Machine*—let alone the recent triumphs of chess-playing programs like Deep Blue, or the World Cup Robot Soccer competitions.

Quite different albeit related considerations could be made in the case of the new 'intelligent' war-machines from the very outset during the First World War. In 1915, for example, the radio-controlled torpedo embodying the simple self-orientation mechanism of Hammond Jr. and Miessner's electric dog was said "to inherit almost superhuman intelligence." The electric dog, initially a "scientific curiosity," Miessner remarked in 1916, "may within the very near future become in truth a real 'dog of war,' without fear, without heart, without the human element so often susceptible to trickery, with but one purpose: to overtake and slay whatever comes within range of its senses at the will of its master." Current development of 'intelligent' weapons make this observation fairly prophetic. And it reminds Wiener's worried judgment on those machines that, as far as their purposeful activity transcend the limitations of the human designer, "may be both effective and dangerous." As Wiener concluded in 1960, the designer could be unable to control a machine whose goal-seeking operation is so fast and irrevocable that, once he become aware that machine's final action has undesired consequences, he had not the time and the data to intervene before such an action is complete.

Various aspects of the discovery of the artificial that are emphasized in the present book have been mostly neglected in investigations on the sciences of the mind and behavior in the twentieth century. A number of authors who play leading roles in this book are not even mentioned or have only a small role in those investigations. Psychologists like Max Meyer or Stephens, let alone engineers or technicians such as Hammond Jr., Miessner, Bent Russell, or Ross, are cases in point. Other authors, who are usually presented as leading figures in the research of the period, are examined in this book in order to bring to light less-known aspects of their research, as in the case of Hull's robot approach or Jennings' views on "inorganic machines." Still other authors, who are towering figures, and justly so, in the usual investigations, are considered here chiefly in order to bring to light internal tensions in their views. The scanty hydraulic analogies of nervous conduction that William James inherited from Herbert Spencer and handed down to William McDougall are an appropriate case in point. Clearly, analogies have an important role in any scientific enterprise. Mechanical analogies for nervous functions, however, occupy a secondary position in the present book, compared to the, albeit naive, *working models* of such functions which were designed or physically realized. Those analogies are discussed in some sections of the present book because examining them allows one to clarify the context in which attempts to build those working models were made; furthermore, they constitute a neglected chapter in the history of psychology and neurology at the threshold of the twentieth century.

Working models or functioning artifacts, rather than scanty analogies, are at the core of the discovery of the artificial, and in two important ways. First, only such models, and not the analogies, can be viewed as tools for testing hypotheses on organism behavior. The way to establish the possibility that complex forms of behavior are not necessarily peculiar to living organisms is "to realize [this possibility] by actual trial," as Hull wrote in 1931 about his own models of learning. And Craik, in *The Nature of Explanation*, contrasted Hull's approach to mechanical models with former, vague "mechanistic views of life and behavior"; in the latter case, "on one hand there has been a tendency to *assert* a mechanistic theory rather than to regard it as a hypothesis which should, if followed out, indicate how and where it breaks down; and on the other hand, there has been little attempt to formulate a definite plan of a mechanism which would fulfill the requirements."

The second way in which working models, rather than analogies, are at the heart of the discovery of the artificial can also be introduced by means of Craik's words: "Any kind of working model of a process is, in a sense, an analogy. Being different it is bound somewhere to break down by showing properties not found in the process it imitates or by not possessing properties possessed by the process it imitates." This is a fundamental statement in the discovery of the artificial: behavioral models themselves can be *tested*. The same point was stated by Grey Walter: "If in the testing [models] fall short of expectation or reality, they do so without equivocation. The model hypothesis cannot bend of flow—it breaks with a loud crack—and from the pieces one can build a better model." And Donald MacKay stated that when a discrepancy between the behavioral model and empirical data is detected, one tries to modify the model in some respect, giving rise to a process of model testing and revision. Behavioral models, Herbert Simon has remarked, "are multicomponent creatures, and when our data don't fit the model, we are faced with a difficult diagnostic task to determine what to change—or whether to discard the entire model." As we shall see, a related and crucial issue is which *constraints*, and how severe, models should satisfy in order to be counted as explanatory tools. In Simon's words, "the important question is not whether we are imitating, but whether in fact the imitation *explains* the phenomenon."

The aim of the present book is not to provide a complete reconstruction of the various chapters in twentieth century neurology, psychology and philosophy of mind, but rather to identify some turning points marking the progressive, albeit discontinuous and tortuous, discovery of the artificial. An effort will be made, however, to present these turning points in their actual contexts, not only to make them more comprehensible, but also to avoid slipping into anachronistic talk of anticipation. Otherwise, the discovery of the artificial would boil down to a catalogue of elementary artifacts and robots: once severed from the context that influenced their construction or evaluation, these would not seem credible bearers of new methodological proposals, distant as they are in time and from current technology.

Even if these pitfalls are avoided, some readers might still remain dumbfounded by the mentalistic interpretations that some designers put on the behavior of their pre-cybernetic machines. One of these machines is said to 'learn,' just because the resistance changes in an electric connection; another is said to 'choose,' just because it automatically opens, or closes, a switch in a circuit. But to use Douglas Hofstadter's words, it is hard to go back and imagine the feeling of those who first saw toothed wheels, hydraulic valves or electrical circuits showing abilities that had been considered exclusive to living organisms. In addition to their being systematically overrated by the public, as already mentioned, simulative methodologies seem bound to meet the following fate: mentalistic descriptions of the artifacts turn out not to be satisfying in the end, and what seems 'intelligent' behavior melts away as soon as it is reproduced in an artifact and becomes, for that very reason, 'automatic.' Behaviorist models seemed trivially simple to those cyberneticians who were aware of them; cybernetic artifacts seemed intrinsically naive to the builders of early intelligent computer programs; and the latter seemed stupid to some of their successors in AI. "AI is whatever hasn't been done yet," was the paradoxical definition reported by Hofstadter.

The discovery of the artificial cannot readily be construed as a cumulative undertaking. At the start, it is the history of isolated efforts, with little contact among its leading promoters. One exception is Hull's robot approach, but that was an experiment that spanned a very short time and whose rapid eclipse, as we shall see, is perhaps exemplary. It is in the 1940s and 1950s that a real turning point comes about in the discovery of the artificial, as a consequence of the rise of computer science and technology, and of automatic control. The diagram below, made by P.L. Simmons and R.F. Simmons in 1960, shows the growth rate in a representative sample of 330 articles and books devoted to behavioral models realized as mathematical simulations, electronic analogs or computer simulations between the end of the 1920s and 1959. In this diagram, the short initial segment of the curve, up to the point corresponding to 1941, would include those isolated attempts at simulative modeling that took place at the outset of the discovery of the artificial, and that are extensively dealt with in the present book (some of them are included in the sample). The publication of Wiener's *Cybernetics* in 1948 coincides with an early peak in the curve, which increases exponentially after 1956, the year of the official birth of AI. In fact, it is beginning with these two dates that the modeling method takes on recognized importance in the community of behavioral scientists, to the point that one begins to speak of the "sciences of the artificial," as Simon put it in the title of his 1969 book. And yet experimentation in the field of the artificial does not become cumulative for this reason. Rather, competing research programmes have coexisted or have

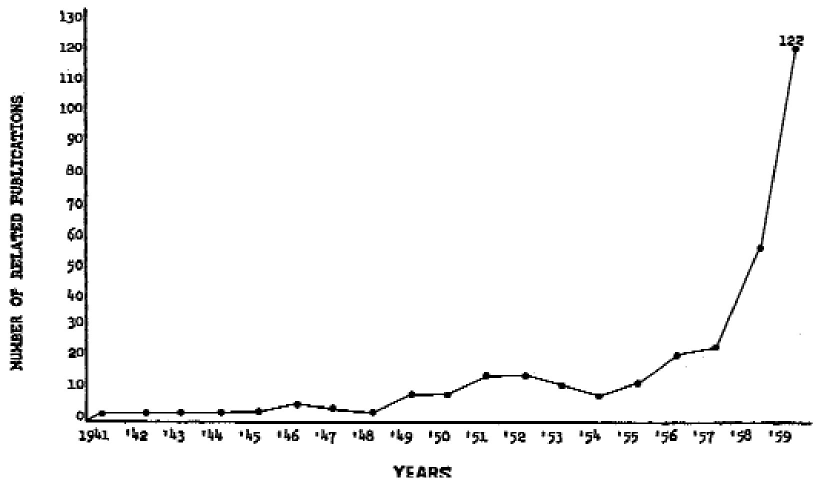

up to now alternated with each other: from different research trends in symbolic AI to current connectionism and neural net approaches and the '*nouvelle*' AI, which includes ALife and different trends in current robotics and synthetic modeling.

To sum up, the discovery of the artificial shares the model-building strategy that is pervasive in other provinces of science. The traditional sciences of the mind—psychology and neurology—addressed during the first half of the twentieth century issues concerning the experimental method and the building and testing of theories of organism behavior and mental life. The earliest behavioral models, designed as simple physical analogs, began to suggest that there might exist a *new* level for testing psychological and neurological hypotheses, one that might coexist alongside the investigations into overt behavior and the nervous system. This was the core idea of what Hull and Craik had already called the "synthetic method," the method of model building. This idea comes fully into focus with the advent of cybernetics, and especially when the pioneers of AI, who wanted to turn their discipline into a new science of the mind, found themselves coming to grips with the traditional sciences of the mind—psychology and neurology—and with their conflicting relationships. That event radically affected the customary taxonomies of the sciences of the mind. From that moment on, the question has explicitly been asked as to what the *right* level of abstraction for constructing explanatory models of mental life might be. Is it the level of 'symbols,' as proposed by classical AI, or that of 'neurons,' as proposed by new connectionism, or those of the genotype and evolutionary and developmental processes, as proposed by ALife, behavior-based and evolutionary robotics and synthetic modeling? And how well-founded are these distinctions?

Although our view of the machine has significantly changed, and is still changing before our eyes, every machine, once taken as a model, is inevitably a simplification of the simulated phenomenon, albeit in relation to the chosen level (or levels) of abstraction. "To what degree is the Rock of Gibraltar a model of the brain?," Ashby once wondered. And he answered, "It persists; so does the brain; they are isomorphic at the lowest level." In fact, the strategy of mental life models faces an unavoidable watershed. It can be maintained that these simplifications do not fatally impoverish the object being studied, once the right levels of abstraction have been found—that they are not "oversimplifications," as Craik put it. Or it can be held that the scientific study of mental life is bound to fail because of the very fact that it makes use of simplifications. In the latter case, hermeneutic or phenomenological intelligibility criteria might be preferred in the study of mind.

I believe that the latter position must face its own difficulties. The final test for models, however, is in the range of phenomena they actually contribute to explaining. The strategy of mental life modeling often raised hopes that it was unable to meet during its history. Nonetheless, this very history provides several lessons for contemporary research, and a wealth of experimental data has been made available that suggests how to establish novel links among the sciences concerned with mental life. The present book tries to document the hopes, defeats and successes in the midst of which the modeling strategy has its origin and growth.

Guides to the discovery of the artificial do exist, and I would like to mention them. Minsky's bibliography in *Computers and Thought*, the classic volume edited in 1963 by Edward Feigenbaum and Julian Feldman, is still extremely useful. The above-mentioned bibliography by P.L. Simmons and R.F. Simmons, "The Simulation of Cognitive Processes. An Annotated Bibliography," in *IRE Transactions on Electronic Computers*, 1961: 462-483; 1962: 535-552, is indispensable. Useful information is contained in various papers and books on cybernetics and the philosophy of psychology in the 1940s and 1950s, such as those by Edwin Boring, F.H. George, T.N. Nemes and Wladislaw Sluckin (these are all included in the list of references, as are other works and authors mentioned above). More than a guide is the set of four volumes edited by Ronald Chrisley, *Artificial Intelligence: Critical Concepts*, Routledge, London and New York, 2000, which collects a number of papers of interest for the discovery of the artificial and its development in the twentieth century.

Unlike other books, such as Martin Gardner's *Logic Machines and Diagrams*, which focuses on logic machines, the present book hardly mentions such machines, except in passing, being more interested in the psychological and neurological use of machines as behavioral models or as tools for understanding mental life. The very selection of the authors discussed here often shows the effects of that interest. As for the decades prior to the cybernetic age, those who tackled the subject of mechanical analogies and "inorganic machines" in the context of the synthetic method were, above all, behaviorist psychologists and neurologists or, before them, scientists and philosophers who had raised the issue of neuro-psychological mechanism.

To facilitate the reading of the book, I have confined to Plates the descriptions of how most machines work. To expedite their interpretation, I have occasionally provided simplified diagrams of some machines. The Plates are not necessary to an understanding of the text, and can be skipped by the reader who is not interested in the details of how those machines work. Most Plates on pre-cybernetic machines contain figures that have never been republished, unless in my earlier articles, or that have never been published at all, like those by Krim.

While writing various essays on the discovery of the artificial during past years, I have had the advantage of discussing the issues addressed in the book with several friends and colleagues, to whom I wish to express my gratitude. Vittorio Somenzi, together with pioneers such as Antonio Borsellino, Valentino Braitenberg, Eduardo Caianiello, Silvio Ceccato and Augusto Gamba, introduced several Italian researchers to cybernetics and AI, and to related philosophical problems in the study of mental life. He has profoundly influenced my training and my research. Giuseppe Trautteur has followed the writing of the book from its very beginning, giving a careful critical reading to its various versions. Ernesto Burattini and Marino Giannini helped me to achieve a better understanding of various points regarding some of the early machine models. Luigia Carlucci Aiello, Massimo Negrotti, Stefano Nolfi, Israel Rosenfield, Guglielmo Tamburrini and Achille Varzi were available on various occasions to read the manuscript and to discuss the issues dealt with in it. Pietro Corsi, Luciano Mecacci and Laurence Smith generously assisted me with their knowledge of the history of psychology and neuroscience. Different chapters of the book were read and usefully commented on by Luciano Floridi, Simone Gozzano, Marcello Frixione, Salvatore Guccione, Domenico Parisi, Silvano Tagliagambe, Settimo Termini and Graziella Tonfoni. Needless to say, none of these is responsible for any error or opinion, even though I have profited from their suggestions and criticism.

Last, but definitely not least, I am indebted to the late Herb Simon for carefully commenting on chapters of the book and for discussing points where our judgement differed. In discussions, he always forced you to go to the heart of a question. His scorn for philosophical disputes not grounded on a “solid body of fact” was only equal to that for idealized models in science that are unable to grasp the complexity of real phenomena.

To write this book has required no small number of visits to libraries around the world. These have been made possible by various grants made available by both the Italian Ministry of University and Scientific and Technological Research (MURST) and the National Research Council (CNR).